arXiv:1610.03803v1 [math.PR] 12 Oct 2016

Applied Probability Trust (13 October 2016)

ROBUST SCHEDULING FOR FLEXIBLE PROCESSING NETWORKS

RAMTIN PEDARSANI,* *University of California, Santa Barbara*

JEAN WALRAND,** *University of California, Berkeley*

YUAN ZHONG,*** *Columbia University*

Abstract

Modern processing networks often consist of heterogeneous servers with widely varying capabilities, and process job flows with complex structure and requirements. A major challenge in designing efficient scheduling policies in these networks is the lack of reliable estimates of system parameters, and an attractive approach for addressing this challenge is to design robust policies, i.e., policies that do not use system parameters such as arrival and/or service rates for making scheduling decisions.

In this paper, we propose a general framework for the design of robust policies. The main technical novelty is the use of a stochastic gradient projection method that reacts to queue-length changes in order to find a balanced allocation of service resources to incoming tasks. We illustrate our approach on two broad classes of processing systems, namely the flexible fork-join networks and the flexible queueing networks, and prove the rate stability of our proposed policies for these networks under non-restrictive assumptions.

Keywords: Robust Scheduling; Flexible Queueing Networks; Stochastic Gradient Projection

* Postal address: Department of Electrical and Computer Engineering, UC Santa Barbara, Santa Barbara, CA 93106, USA. Email: ramtin@ece.ucsb.edu

** Postal address: 257 Cory Hall, Department of Electrical Engineering and Computer Sciences, UC Berkeley, Berkeley, CA 94720, USA. Email: walrand@berkeley.edu

*** Postal address: 500 W. 120th St., Mudd 344, New York, NY 10027, USA. Email: yz2561@columbia.edu

2010 Mathematics Subject Classification: Primary 60K25
Secondary 90B15; 60G17

1. Introduction

As modern processing systems (e.g., data centers, hospitals, manufacturing networks) grow in size and sophistication, their infrastructures become more complicated, and a key operational challenge in many such systems is the efficient scheduling of processing resources to meet various demands in a timely fashion. A scheduling policy decides how server capacities are allocated over time, and a major challenge in designing such policies is the lack of knowledge of system parameters due to the complex processing environment and the volatility of the jobs to be processed. Demands are diverse, typically unpredictable, and can occur in bursts; furthermore, operating conditions of processing resources can vary over time (see e.g., [18]). Thus, estimates of key system parameters such as arrival and/or service rates are often unreliable, and can frequently become obsolete. One approach to address this complicated scheduling and resource allocation problem is to design *robust* scheduling policies, where scheduling decisions are made only based on current and/or past system states such as queue sizes, and not depending on system parameters such as arrival or service rates. Robust scheduling policies can be highly desirable in practice, since they use only minimal information and can adapt to changes in demands and service conditions automatically. The main objective of this paper is to develop a general framework for designing robust policies and analyzing their performance.

Consider a single-server queueing system with unit-size jobs arriving at an unknown rate λ, and a server with a costly and sufficiently large service capacity μ (in particular, $\mu > \lambda$), whose precise value is unknown. Suppose that at regular time intervals, the server can adjust its service effort, measured by the fraction $p \in [0, 1]$ of the total capacity (we can implement this in practice as a randomized decision of serving the queue with probability p). The goal is to keep the system stable. Let ΔQ be the queue size change over a regular time interval. Intuitively, if $\Delta Q > 0$, then it is likely that the arrival rate is faster than the dedicated service effort, which should then be increased. If $\Delta Q < 0$, then the service effort should be decreased for cost consideration. This naturally leads to an update rule for the service effort from time n to $n+1$ of the form

$p^{n+1} := p^n + \beta^n \Delta Q$ with $\beta^n > 0$. Under mild technical conditions on the sequence $\{\beta^n\}$, it can be shown that $\mu^n \to \lambda$ almost surely, implying system stability.

This simple example illustrates the high-level approach of our policy design framework: allocate more (less) service to a queue if the corresponding queue size increases (decreases). Our design approach only uses the system state information – namely, the queue size changes – and does not require information on either the arrival or service rates. A simple but key observation that justifiess the validity of this approach is that if the queue size at the start of an interval is sufficiently large, then ΔQ, the queue size change, is proportional to $\lambda - \mu p$ in expectation. In a network setting, by building upon this simple observation, we can decide how to allocate shared resources among competing queues based on their respective queue size changes.

Our methodology is general and can be applied to a wide range of processing networks. To illustrate our approach concretely, we focus on two broad classes of network models, which (a) generalize many important classes of queueing network models, such as parallel server systems [20] and fork-join networks [21] (see Section 1.1 for more details), and (b) model key features of dynamic resource allocation at fine granularity in many modern applications such as cloud computing, flexible manufacturing, and large-scale healthcare systems. We now proceed to describe our network models and contributions in more detail.

Common to many modern large-scale processing systems are the following two important features: (a) workflows of interdependent tasks, where the completion of one task produces new tasks to be processed in the system, and b) flexibility of processing resources with overlapping capabilities as well as flexibility of tasks to be processed by multiple servers. To illustrate these two features, consider the scheduling of a simple Mapreduce job [12] of word count of the play "Hamlet" in a data center (see Figure 1). *"Mappers"* are assigned the tasks of word count by Act, producing intermediate results, which are then aggregated by the "reducer". In more elaborate workflows, these interdependencies can be more complicated. There may also be considerable overlap in the processing capabilities of the data center servers, and flexibility on where tasks can be placed [12]. Similarly, in a healthcare facility such as a hospital, an arriving patient may have a complicated workflow of service/treatment requirements [2], which can also be assigned to doctors and/or nurses with overlapping capabilities. To capture

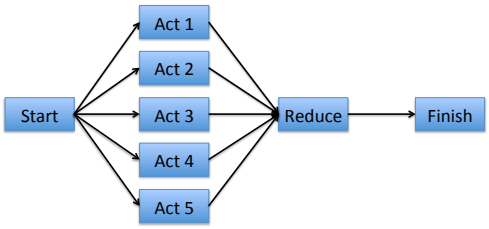

FIGURE 1: Word count of Hamlet in MapReduce

the dependencies in workflows and the system flexibility, we consider the following two classes of processing networks:

(i) a flexible fork-join processing network model, in which jobs are modeled as directed acyclic graphs (DAG), with nodes representing *tasks*, and edges representing *precedence constraints* among tasks, and servers have overlapping capabilities; and

(ii) a flexible queueing network with probabilistic routing structure, where a job goes through processing steps in different queues according to a routing matrix, and servers have overlapping capabilities.

We design a robust scheduling policy for each class of networks, and analyze performance properties of the proposed policies. For both models, we prove the throughput optimality[1] of our policies, under a factorization criterion on service rates (see Assumption 1 of Section 2.4 for details). Our policy design is based on the simple idea of matching incoming flow rates to their respective service rates, and detecting mismatches using queue size information. If system parameters were known, a so-called static planning problem [15] can be solved to obtain the optimal allocation of server capacities, which balances flows in the system. Without the knowledge of system parameters, however, the policy updates the allocation of server capacities according to changes in queue sizes. Methodologically, our policy uses the idea of stochastic gradient descent (see e.g., [7]), a technique that has been applied in the design of distributed policies in ad-hoc wireless networks [17].

[1]We are concerned with *rate stability* in this paper. A scheduling policy is *throughput optimal* if, under this policy, the system is stable whenever there exists some policy under which the system is stable.

1.1. Related Works

Scheduling of queueing networks has been studied extensively over several decades. We do not attempt to provide a comprehensive literature review here; instead we review the most relevant works.

Our flexible fork-join network model is closely related to and substantially generalizes the classical fork-join networks (see e.g., [4, 3, 19, 6, 21, 22]). The main difference between the classical models and ours is that we allow tasks to be flexible, whereas tasks are assigned to dedicated servers in classical fork-join networks. In classical networks, simple robust policies such as FIFO (First-In-First-Out) can often shown to be throughput optimal, but need not be so in our flexible networks.

The flexible queueing network model is closely related to the system considered in [1] ([1] also considers setup costs whereas we do not). The policies in [1] make use of arrival and service rates, and their throughput properties are analyzed using fluid models, hence their approach is distinct from ours. We would also like to point out that in the case where the queues are not flexible, i.e., each queue has a dedicated server, the system reduces to the open multiclass queueing network (see e.g., [16, 11]).

Both the flexible fork-join network model and the flexible queueing network model can be viewed as generalizations of the classical parallel server system, considered in e.g., [20]. The flexible fork-join network extends the parallel server system by allowing jobs to consist of tasks with precedence constraints, and the flexible queueing network extends the parallel server system by allowing probabilistic routing among jobs. There is considerable interest in the study of robust scheduling algorithms in the context of parallel server systems. The well-know $Gc\mu$ rule – equivalent to a MaxWeight policy with appropriately chosen weights on queues – has been proved to have good performance properties, including throughput optimality (e.g., [20]). The $Gc\mu$ rule does not make use of the knowledge of arrival rates, but does require the knowledge of service rates. [5] studies performance properties of the Longest-Queue-First (LQF) policy, which is robust to both arrival and service rates, and establishes its throughput optimality when the so-called activity graph is a tree. [26] established the throughput optimality of a priority discipline in a many-server parallel server system, which consists of server pools, each of which in turn consists of a large number of identical servers, also

under the condition that the activity graph is a tree. [13] established the throughput optimality of LQF under a local pooling condition. To the best of our knowledge, in the flexible fork-join networks and the flexible queueing networks, both extensions of parallel server systems, which include precedence constraints and routing, respectively, the problem of designing robust scheduling policies has not been addressed prior to this work.

As mentioned earlier, the analysis of our policies uses the technique of stochastic gradient descent [7], which has been successfully employed in the design of distributed CSMA algorithms for wireless networks [17]. Our analysis is different from that of the CSMA algorithms in several ways; for example, CSMA algorithms actively attempt to estimate the arrival and service rates, whereas our policy is adaptive, and only reacts to these parameters through queue size changes.

1.2. Paper Organization

The rest of the paper is organized as follows. In Section 2, we introduce the flexible fork-join network model. We propose our robust scheduling policy, and state our main theorems. In Section 3, we describe the flexible queueing network model, and design a robust scheduling policy for this network. We conclude the paper in Section 4. All proofs are provided in the appendices.

2. Scheduling DAGs with Flexible Servers

2.1. System Model

We consider a general flexible fork-join processing network, in which jobs are modeled as directed acyclic graphs (DAG). Jobs arrive to the system as a set of tasks, among which there are precedence constraints. Each node of the DAG represents one task type[2], and each (directed) edge of the DAG represents a precedence constraint. More specifically, we consider M classes of jobs, each of them represented by one DAG structure. Let $\mathcal{G}_m = (\mathcal{V}_m, \mathcal{E}_m)$ be the graph representing the job of class m, $1 \leq m \leq M$, where $\mathcal{V}_m$ denotes the set of nodes of type-m jobs, and $\mathcal{E}_m$ the set of

[2]We make use of both the concepts of *tasks* and *task types*. To avoid confusion and overburdening terminology, we will often use *node* synonymously with *task type* for the rest of the paper.

edges of the graph. Note that sets $\mathcal{V}_m$ are disjoint. Let $\mathcal{V} = \cup_{m=1}^{M} \mathcal{V}_m$ and $\mathcal{E} = \cup_{m=1}^{M} \mathcal{E}_m$. We suppose that each $\mathcal{G}_m$ is connected, so that there is an undirected path between any two nodes of $\mathcal{G}_m$. There is no directed cycle in any $\mathcal{G}_m$ by the definition of DAG. Let the number of nodes of job type m be K_m, i.e. $|\mathcal{V}_m| = K_m$. Let the total number of nodes in the network be K. Thus, $\sum_{m=1}^{M} K_m = K$. We index the task types in the system by k, $1 \leq k \leq K$, starting from job type 1 to M. Thus, task type k belongs to job type $m(k)$ if

$$\sum_{m'=1}^{m(k)-1} K_{m'} < k \leq \sum_{m'=1}^{m(k)} K_{m'}.$$

We call node k' a *parent* of node k, if they belong to same job type m, and $(k', k) \in \mathcal{E}_m$. Let $\mathcal{P}_k$ denote the set of parents of node k. In order to start processing a type-k task, the processing of all tasks of its parents *within the same job* should be completed. Node k is said to be a *root* of DAG type $m(k)$, if $\mathcal{P}_k = \emptyset$. We call k' an *ancestor* of k if they belong to the same DAG, and there exists a directed path of edges from k' to k. Let L_k be the length of the longest path from the root nodes of the DAG, $\mathcal{G}_{m(k)}$, to node k. If k is a root node, then $L_k = 0$.

There are J servers in the processing network. A server is *flexible* if it can serve more than one type of tasks. A task type is flexible if it can be served by more than one server. In other words, servers can have overlap of capabilities in processing a node. For each j, we define $\mathcal{T}_j$ to be the set of nodes that server j is capable of serving. Let $T_j = |\mathcal{T}_j|$. For each k, let $\mathcal{S}_k$ be the set of servers that can serve node k, and let $S_k = |\mathcal{S}_k|$. Without loss of generality, we also assume that $T_j, S_k \geq 1$ for all j and k, so that each server can serve at least one node, and each node can be served by at least one server.

Example 1. Figure 2a illustrates the DAG of one job type that consists of four nodes $\{1, 2, 3, 4\}$. There are two servers 1 and 2. Server 1 can process tasks of types in the set $\mathcal{T}_1 = \{1, 2, 3\}$ and server 2 can process tasks of types in the set $\mathcal{T}_2 = \{3, 4\}$. When a type-1 task is completed, it "produces" one type-2 task and one type-3 task, both of which have to be completed before the processing of the type-4 task of the same job can start.

We consider the system in discrete time. We assume that the arrival process of

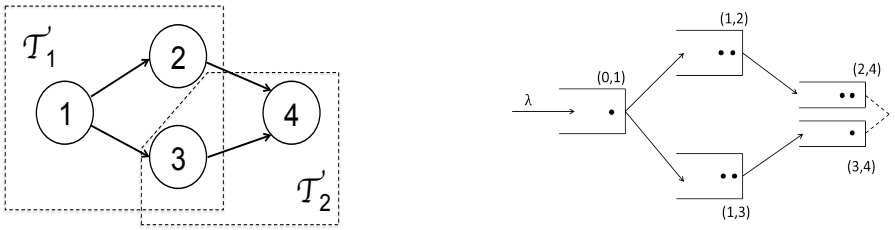

(a) A simple DAG (b) Queueing Network of the DAG

FIGURE 2: This figure illustrates a simple DAG and its corresponding queueing network.

type-m jobs is an independent Bernoulli process with rate λ_m, $0 < \lambda_m < 1$; that is, in each time slot, a new job of type m arrives to the system with probability λ_m, independently over time. We assume that the service times are geometrically distributed and independent of everything else. When server j processes task k, the service completion time has mean μ_{kj}^{-1}. Thus, μ_{kj} can be interpreted as the service rate of node k when processed by server j.

2.2. Queueing Network Model for Cooperative Servers

We model our processing system as a queueing network in the following manner. We maintain one virtual queue of processed tasks that are sent from node k' to k for each edge of the DAGs $(k', k) \in \mathcal{E}$. Furthermore, we maintain a virtual queue for the root nodes of the DAGs. Let χ_m be the number of root nodes in the graph of job type m. Then, the queueing network has $\sum_{m=1}^{M}(|\mathcal{E}_m| + \chi_m)$ virtual queues. As an example, consider the DAG of Figure 2a. The virtual queues corresponding to this DAG is illustrated in Figure 2b. There are 5 virtual queues in total, 4 for edges of the graph, and one for the root node 1.

Job identities and synchronization. In our model, jobs and tasks have distinct identities. This is mainly motivated by applications to data centers and healthcare systems. For instance, it is important not to mix up blood samples of different patients in hospital, and to put pictures on the correct webpage in a data center setting.

To illustrate how a job is processed in the preceding queueing network, consider the network in Figure 2b. Suppose that a job of identity a and with a DAG structure of Figure 2a enters the network. When task 1 of job a from queue $(0, 1)$ is processed,

tasks 2 and 3 of job a are sent to queues $(1,2)$ and $(1,3)$, respectively. When tasks in queues $(1,2)$ and $(1,3)$ are processed, their results are sent to queues $(2,4)$ and $(3,4)$, respectively. Finally, to process task 4 of job a, one task belonging to job a from queue $(2,4)$ and one task belonging to job a from $(3,4)$ are gathered and processed to finish processing job a. We emphasize that tasks are identity-aware in the sense that to complete processing task 4, it is not possible to merge any two tasks (of possibly different jobs) from queues $(2,4)$ and $(3,4)$.

A common and important problem that needs to be addressed in scheduling DAGs is *synchronization*, where all parents of a task need to be completed for the task to be processed. In the presence of flexibility, synchronization constraints may lead to disorder in task processing, which adds synchronization penalty to the system; see [24] for an example. In this paper, to guarantee synchronization, we assume the simplifying condition that servers are cooperative (this is equivalent to the case of cooperating servers described in [1]). That is, we assume that servers that work on the same task type, cooperate on the same head-of-the-line task, adding their service capacities. In this way, tasks are processed in a FIFO manner so that no synchronization penalty is incurred.

Queue Dynamics. Now we describe the dynamics of the queueing network. Let $Q_{(k',k)}$ denote the length of the queue corresponding to edge (k',k) and let $Q_{(0,k)}$ denote the length of the queue corresponding to root node k. A task of type k can be processed if and only if $Q_{(k',k)} > 0$ for all $k' \in \mathcal{P}_k$ – this is because servers are cooperative, and tasks are processed in a FIFO manner. Thus, the number of tasks of node k available to be processed is $\min_{k' \in \mathcal{P}_k} Q_{(k',k)}$, if k is not a root node, and $Q_{(0,k)}$, if k is a root node. For example, in Figure 2b, queue $(2,4)$ has length 2 and queue $(3,4)$ has length 1, so there is one task of type 4 available for processing. When one task of class k is processed, lengths of all queues (k',k) are decreased by 1, where $k' \in \mathcal{P}_k$, and lengths of all queues (k,i) are increased by 1, where $k \in \mathcal{P}_i$. Therefore, the dynamics of the queueing network is as follows. Let $d_k^n \in \{0,1\}$ be the number of processed tasks of type k at time n, and $a_m^n \in \{0,1\}$ be the number of jobs of type m

that arrives at time n. If k is a root node of the DAG, then

$$Q^{n+1}_{(0,k)} = Q^n_{(0,k)} + a^n_{m(k)} - d^n_k; \tag{1}$$

else,

$$Q^{n+1}_{(k',k)} = Q^n_{(k',k)} + d^n_{k'} - d^n_k. \tag{2}$$

Let p_{kj} be the fraction of capacity that server j allocates for processing available tasks of class k. We define $p = [p_{kj}]$ to be the *allocation vector*. If server j allocates all its capacity to different tasks, then $\sum_{k\in\mathcal{T}_j} p_{kj} = 1$. Thus, an allocation vector p is called *feasible* if

$$\sum_{k\in\mathcal{T}_j} p_{kj} \leq 1, \ \forall\ 1 \leq j \leq J. \tag{3}$$

We interpret the allocation vector at time n, $p^n = [p^n_{kj}]$, as randomized scheduling decisions at time n, in the following manner. First, without loss of generality, the system parameters can always be re-scaled so that $\sum_{j\in\mathcal{S}_k} \mu_{kj} < 1$ for all k, by speeding up the clock of the system. Now suppose that at time slot n, the allocation vector is p^n. Then, the head-of-the-line task k is served with probability $\sum_{j\in\mathcal{S}_k} \mu_{kj} p^n_{kj}$ in that time slot. Note that $\sum_{j\in\mathcal{S}_k} \mu_{kj} p^n_{kj} < 1$ by our scaling of the service rates.

2.3. The Static Planning Problem

In this subsection, we introduce a linear program (LP) that characterizes the *capacity region* of the network, defined to be the set of all arrival rate vectors λ where there is a scheduling policy under which the queueing network of the system is stable[3]. The *nominal* traffic rate to all nodes of job type m in the network is λ_m. Let $\nu = [\nu_k] \in \mathbb{R}^K_+$ be the set of nominal traffic rate of nodes in the network. Then, $\nu_k = \lambda_m$ if $m(k) = m$, i.e., if $\sum_{m'=1}^{m-1} K_{m'} < k \leq \sum_{m'=1}^{m} K_{m'}$. The LP that characterizes the capacity region of the network makes sure that the total service capacity allocated to each node in the network is at least as large as the nominal traffic rate to that node. Formally, the LP – known as the *static planning problem* [15] – is defined as follows.

[3]As mentioned earlier, the stability condition that we are interested in is rate stability.

$$\begin{aligned}
\text{Minimize} \quad & \rho && (4)\\
\text{subject to} \quad & \nu_k \le \sum_{j\in\mathcal{S}_k} \mu_{kj}p_{kj},\ \forall\ 1\le k\le K\\
& \rho \ge \sum_{k\in\mathcal{T}_j} p_{kj}, \qquad \forall\ 1\le j\le J, && (5)\\
& p_{kj}=0, \qquad \text{if } k\notin\mathcal{T}_j, && (6)\\
& p_{kj}\ge 0. && (7)
\end{aligned}$$

Proposition 2.1. *Let the optimal value of the LP be ρ^*. Then $\rho^* \le 1$ is a necessary and sufficient condition of rate stability of the system under some scheduling policy.*

The proof of Proposition 2.1 is provided in Appendix 5.1.

Thus, by Proposition 2.1, the *capacity region* Λ of the network is the set of all $\lambda \in \mathbb{R}_+^M$ for which the corresponding optimal solution ρ^* to the LP satisfies $\rho^* \le 1$. More formally,

$$\Lambda \triangleq \left\{ \lambda \in \mathbb{R}_+^M : \exists\ p_{kj}\ge 0 \text{ such that } \sum_{k\in\mathcal{T}_j} p_{kj} \le 1\ \ \forall\ j, \text{ and } \nu_k \le \sum_{j\in\mathcal{S}_k} \mu_{kj}p_{kj}\ \ \forall\ k \right\}.$$

2.4. Scheduling Policy Robust to Task Service Rates

In this subsection, we make the following assumption on service rates μ_{kj}.

Assumption 1. *For all k and all $j \in \mathcal{S}_k$, service rates μ_{kj} can be factorized to two terms: a task-dependent term μ_k, and a server-dependent term α_j. Thus, $\mu_{kj} = \mu_k\alpha_j$.*

While Assumption 1 appears somewhat restrictive, it covers a variety of important cases. When $\alpha_j = 1$ for all j, the service rates are task dependent. This case models, for example, a data center of servers with the same processing speed (possibly of the same generation and purchased from the same company), but with different software compatibilities, and possibly hosting overlapping sets of data blocks. The case when α_j are different can model the inherent heterogeneous processing speeds of the servers.

We now propose a scheduling policy with known α_j, which is robust to task service rates μ_k, and prove that it is throughput optimal. The idea of our scheduling policy is quite simple: it reacts to queue size changes by adjusting the service allocation vector $p = [p_{kj}]$. Since service rates μ_{kj} are factorized to two terms μ_k and α_j, only the sum

$p_k \triangleq \sum_{j\in\mathcal{S}_k} \alpha_j p_{kj}$ affects the effective service rate for node k. One can consider p_k as the total capacity that all the servers allocate to node k in a time slot. So, with a slight abuse of notation and terminology, we call $p = [p_k]$ the service allocation vector from now on.

To precisely describe our proposed scheduling algorithm, first we introduce some notation. Let $\mathbf{1}\{Q^n_{(k',k)} > 0\}$ be the indicator that the queue corresponding to edge (k',k) is non-empty at time n. Let $\Delta Q^{n+1}_{(k',k)} = Q^{n+1}_{(k',k)} - Q^n_{(k',k)}$ be the size change of queue (k',k) from time n to $n+1$. Define E^n_k to be the event that there is a strictly positive number of type-k tasks to be processed at time n. Thus, $E^n_k = \{Q^n_{(0,k)} > 0\}$ if k is a root node, and $E^n_k = \{Q^n_{(k',k)} > 0,\ \forall k' \in \mathcal{P}_k\}$ if k is not a root node. Also let $\mathbf{1}_{E^n_k}$ be the indicator function of event E^n_k.

Let $\mathcal{C} \subseteq \mathbb{R}^K_+$ be the polyhedron of feasible service allocation vector p.

$$\mathcal{C} = \left\{ p \in \mathbb{R}^K : \exists\ p_{kj} \text{ such that } \sum_{j\in\mathcal{S}_k} \alpha_j p_{kj} = p_k \ \ \forall k,\ p_{kj} \geq 0 \ \ \forall k,j,\ \sum_{k\in\mathcal{T}_j} p_{kj} \leq 1 \ \ \forall j \right\}. \tag{8}$$

For any K-dimensional vector x, let $[x]_{\mathcal{C}}$ denote the convex projection of x onto $\mathcal{C}$. Finally, let $\{\beta^n\}$ be a positive decreasing sequence with the following properties: (i) $\sum_{n=1}^{\infty} \beta^n = \infty$, (ii) $\sum_{n=1}^{\infty} (\beta^n)^2 < \infty$, and (iii) $\lim_{n\to\infty} \frac{1}{n\beta^n} < \infty$.

As we will see in the sequel, a key step of our algorithm is to find an unbiased estimator of $\lambda - \mu_k p_k$ for all k, based on the current and past queue sizes. Toward this end, for each node k, we first pick a path of queues from a queue corresponding to a root node of the DAG to queue (k',k) for some $k' \in \mathcal{P}_k$. Note that the choice of this path need not be unique. Let $\mathcal{H}_k$ denote the set of queues on this path from a root node to node k. For example, in the DAG of Figure 2a, for node 4, we can pick the path $\mathcal{H}_4 = \{(0,1),(1,2),(2,4)\}$. Then, we use $\sum_{(i',i)\in\mathcal{H}_k} \Delta Q_{(i',i)}$ as an unbiased estimate of $\lambda - \mu_k p_k$. To illustrate the reason behind this estimate, consider the DAG in Figure 2b. It is easy to see that

$$\mathbb{E}(\Delta Q^{n+1}_{(0,1)} + \Delta Q^{n+1}_{(1,2)} + \Delta Q^{n+1}_{(2,4)} | Q^n_{(2,4)} > 0, Q^n_{(3,4)} > 0, p^n) = \lambda - \mu_4 p^n_4.$$

In general, if $L_k = L - 1$ for node k (recall that L_k is the length of the longest path from a root node to k), one picks a path of edges $(i_0,i_1),(i_1,i_2),\ldots,(i_{L-1},i_L)$, such

that $i_0 = 0$ and $i_L = k$. Then,

$$\mathbb{E}\left[\sum_{l=0}^{L-1} \Delta Q^{n+1}_{(i_l,i_{l+1})} | \mathbf{1}_{E^n_k} = 1, p^n\right] = (\nu_k - \mu_{i_1} p^n_{i_1} \mathbf{1}_{E^n_{i_1}}) + \sum_{l=1}^{L-1} (\mu_{i_l} p^n_{i_l} \mathbf{1}_{E^n_{i_l}} - \mu_{i_{l+1}} p^n_{i_{l+1}} \mathbf{1}_{E^n_{i_{l+1}}})$$
$$= \nu_k - \mu_k p^n_k \mathbf{1}_{E^n_k}. \tag{9}$$

Note that one can pick any path from a root node to k, but the longest path is picked in (9) for the purpose of ease of notation for the proofs. Our scheduling algorithm updates the allocation vector p^n in each time slot n in the following manner.

1. We initialize with an arbitrary feasible p^0.
2. Update the allocation vector p^n as follows:

$$p^{n+1}_k = [p^n_k + \beta^n \mathbf{1}_{E^n_k} \sum_{(i',i) \in \mathcal{H}_k} \Delta Q^{n+1}_{(i',i)}]_{\mathcal{C}}. \tag{10}$$

This completes the description of the algorithm.

We now provide some intuition for the algorithm. As we mentioned, the algorithm tries to find adaptively the capacity allocated to task k, p_k, that balances the nominal arrival rate and departure rate of queues (k', k). The nominal traffic of all the queues of DAG type m is ν_k for task types k belonging to job type m. Thus, the algorithm tries to find $p^*_k = \frac{\nu_k}{\mu_k}$, in which case the nominal service rate of all the queues is $p^*_k \mu_k = \nu_k$. To find an adaptive robust algorithm, we formulate the following optimization problem.

$$\begin{aligned} \text{minimize} \quad & \frac{1}{2} \sum_{k=1}^{K} (\nu_k - \mu_k p_k)^2 \\ \text{subject to} \quad & p \in \mathcal{C}. \end{aligned} \tag{11}$$

Solving (11) by the standard gradient descent algorithm, using step size β^n at time n, leads to the update rule

$$p^{n+1}_k = [p^n_k + \beta^n \mu_k (\nu_k - \mu_k p^n_k)]_{\mathcal{C}}. \tag{12}$$

To make the update in (12) robust, first we consider a "skewed" update

$$p^{n+1}_k = [p^n_k + \beta^n (\nu_k - \mu_k p^n_k)]_{\mathcal{C}}, \tag{13}$$

and second, we use the queue-length changes in (9) as an unbiased estimator of the term $\nu_k - \mu_k p^n_k$. This results in the update equation in (10). Thus, the update in

(10) becomes robust to knowledge of service rates μ_k. The algorithm is not robust to knowledge of server rates, α_j, since the convex set $\mathcal{C}$ is dependent on α_j, and the projection requires the knowledge of server speeds α_j.

Let us now provide some remarks on the implementation efficiency of the algorithm. First, the policy is not fully distributed. While the update variables $\mathbf{1}_{E_k^n} \sum_{(i',i)\in\mathcal{H}_k} \Delta Q_{(i',i)}^{n+1}$ can be computed locally, the projection $[\cdot]_{\mathcal{C}}$ requires full knowledge of all these local updates. Second, since Euclidean projection on a polyhedron is a quadratic programming problem that can be solved efficiently in polynomial time by optimization algorithms such as the "interior point method" [8], the projection step $[\cdot]_{\mathcal{C}}$ can be implemented efficiently.

The simulation results are derived using the described algorithm. To analyze the theoretical performance properties of the algorithm, we make minimal modifications to the proposed algorithm for technical reasons. First, we assume that a) the nominal arrival rate of all the tasks ν_k is strictly positive, b) there are finitely many servers in the system, and c) all the service rates, μ_{kj}, are finite. Note that assumptions a), b), and c) can be made without loss of generality. Then, there exists $\varepsilon_0 > 0$ such that for all k, $\frac{\nu_k}{\mu_k} \geq \varepsilon_0$. We now suppose that ε_0 is known, and consider a variant $\mathcal{C}_{\varepsilon_0}$ of the convex set $\mathcal{C}$, defined to be

$$\mathcal{C}_{\varepsilon_0} = \left\{ p \in \mathbb{R}^K : \exists\ p_{kj} \geq 0 \text{ such that } \sum_{j\in\mathcal{S}_k} \alpha_j p_{kj} = p_k \ \ \forall k,\ p_k \geq \varepsilon_0 \ \ \forall k,\ \sum_{k\in\mathcal{T}_j} p_{kj} \leq 1 \ \ \forall j \right\}. \tag{14}$$

Note that $p^* \in \mathcal{C}_{\varepsilon_0}$. We modify the projection to be on the set $\mathcal{C}_{\varepsilon_0}$ every time, so that p^n are now updated as

$$p_k^{n+1} = [p_k^n + \beta^n \mathbf{1}_{E_k^n} \sum_{(i',i)\in\mathcal{H}_k} \Delta Q_{(i',i)}^{n+1}]_{\mathcal{C}_{\varepsilon_0}}. \tag{15}$$

Remark 1. a) The modified update (15) ensures that the service effort $\mu_k p_k^n$ allocated to each queue k is always strictly between 0 and 1, which will imply that all queues are non-empty for a positive fraction of the time, so as to guarantee convergence of the modified algorithm (15). We believe that the original algorithm (10) converges as well, although establishing this fact rigorously appears difficult.

b) Let us also note that the modified algorithm (15) is *essentially* robust in the following sense. On the one hand, the update (15) assumes the knowledge of ε_0, which in turn depends on ν and μ. On the other hand, ε_0 can be chosen with minimal information on ν and μ. For example, if c is a known lower bound on $\nu_k, k = 1, 2, \cdots, K$ and C is a known upper bound on $\mu_k, k = 1, 2, \cdots, K$, then we can set $\varepsilon_0 = \frac{c}{C}$.

The main results of this section are the following two theorems.

Theorem 2.1. *Let $\lambda \in \Lambda$. The allocation vector p^n updated by Equation (15) converges to $p^* = [p_k^*]$ almost surely, where $p_k^* = \frac{\nu_k}{\mu_k}$.*

The proof of Theorem 2.1 is provided in Appendix 5.2. Here, we briefly describe the main steps of the proof. First, we show that the non-stochastic gradient projection algorithm with the skewed update (13) converges. This is not true in general, but the convergence holds here, due to the form of the objective function in (11), which is the sum of separable quadratic terms. Second, we show that the cumulative stochastic noise present in the update due to the error in estimating the correct drift is an L_2-bounded martingale. Thus, by martingale convergence theorem the cumulative noise converges and has a vanishing tail. This shows that after some time the noise becomes negligible. Finally, we prove that the event that all queues in the network are non-empty happens for a positive fraction of time. Intuitively, this suggests that the algorithm is updating "often enough" to be able to converge. The rigorous justification makes use of Kronecker's lemma [14].

Theorem 2.2. *Let $\lambda \in \Lambda$. The queueing network representing the DAGs is rate stable under the proposed scheduling policy, i.e. $\lim_{n\to\infty} \frac{Q^n_{(k',k)}}{n} = 0,\ a.s.,\ \forall (k', k)$.*

The proof of Theorem 2.2 is provided in Appendix 5.3. While proving the theorem is technically quite involved, the key idea is to use Theorem 2.1 to prove that the servers allocate enough cumulative capacity to all the tasks in the system, which leads to rate stability of the network.

2.5. Simulations

In this section, we show the simulation results and discuss the performance of the robust scheduling algorithm. Consider the DAG shown in Figure 3. We assume that

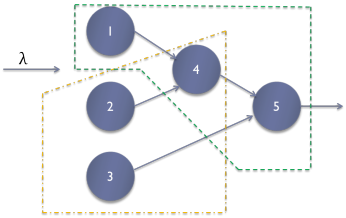

FIGURE 3: DAG of 5 tasks

the system has 1 type of jobs with arrival rate $\lambda = 0.23$. The task service rates are

$$\mu_1 = 1, \mu_2 = 4/3, \mu_3 = 2, \mu_4 = 1/2 \text{ and } \mu_5 = 2/3.$$

The server speeds are $\alpha_1 = 1$ and $\alpha_2 = 1/2$, and $\mathcal{T}_1 = \{1, 4, 5\}$ and $\mathcal{T}_2 = \{2, 3, 4\}$. The step size of the algorithm is chosen to be $\beta^n = \frac{1}{n^{0.6}}$ and the initial queue lengths are $[0, 0, 0]$. From (4), one can compute that the capacity region Λ is $\{\lambda \geq 0 : \lambda \leq \frac{6}{23}\}$.

First we demonstrate that our proposed algorithm is throughput optimal and makes the queues stable. Figure 4a illustrates the queue-lengths as a function of time. Moreover, Figure 4b shows how vector p^n converges to the flow-balancing values as Theorem 2.1 states. Figure 4a suggests that queues in the network become empty infinitely often, hence are stable. However, the average queue-length is quite large, so the algorithm suffers from bad delay. The reason is that the allocation vector p^n is converging to the value that equalizes the arrival and service rates of all the queues. As an example, if we consider a system with a single-node DAG of arrival rate λ and a single server of service rate μ, the queueing network reduces to the classical M/M/1 queue. Theorem 2.1 shows that the service capacity that this queue receives, p^n, converges to $\frac{\lambda}{\mu}$. It is known that if the arrival rate of an M/M/1 queue is equal to its service rate, the underlying Markov chain describing the queue-length evolution is null-recurrent, and the queue suffers from large delay. In the following, we propose a modified version of the algorithm that reduces the delays.

Modified scheduling algorithm to improve delays. As discussed in the previous section, the allocation vector p^n converges to the value that just equalizes the arrival rate and the effective service rate that each tasks receives. To improve the delay of the system, one wants to allocate strictly larger service rate to each task than the arrival

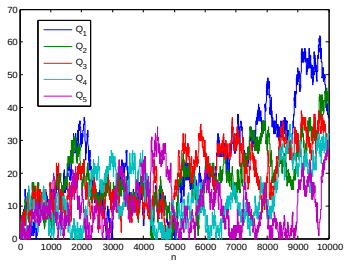

(a) Queue-lengths vs. time

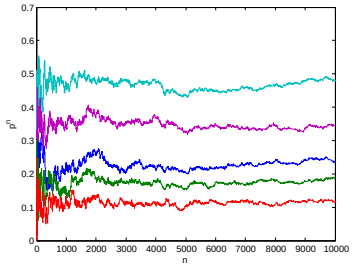

(b) Allocation vector p^n vs. time

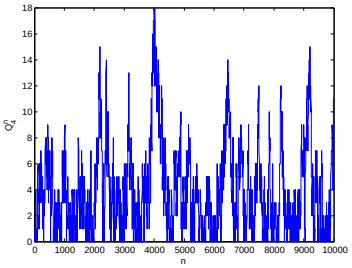

(c) Queue-length of queue 4 for the modified algorithm vs. time for $\delta = 0.02$.

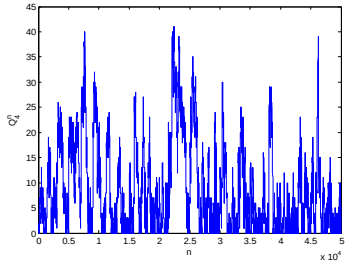

(d) Queue-length evolution of queue 4 in the time-varying bursty case.

FIGURE 4: This figure shows the simulation results for the DAG of Figure 3.

rate. This is possible only if the arrival vector λ is in the interior of the capacity region. In this case, there exists some $\delta > 0$ and an allocation vector p^* such that $\nu_k \leq -\delta + \mu_k p_k^*$ for all k.

Thus, assuming that δ is known, we minimize the function

$$V(p) = \sum_{k=1}^{K} (\nu_k + \delta - \mu_k p_k^*)^2,$$

by stochastic gradient. Similarly to the proof of Theorem 2.1, one can show that p_k^n converges to $\frac{\nu_k + \delta}{\mu_k}$. With this formulation of the optimization problem, the update equation for the new scheduling algorithm is

$$p_k^{n+1} = [p_k^n + \delta + \beta^n \mathbf{1}_{E_k^n} \sum_{(i',i) \in \mathcal{H}_k} \Delta Q_{(i',i)}^{n+1}]_{\mathcal{C}},$$

which is similar to (10) with an extra δ-slack.

We consider the same setting and network parameters as the previous section, and simulate the modified algorithm using $\delta = 0.02$. Figure 4c demonstrates a substantial

reduction in the queue-length of queue 4 and the delay performance of the algorithm. Similar plots can be obtained for queue-lengths of other queues, which we omit to avoid redundancy.

Time-varying demand and service and bursty arrivals. In Section 1, we mentioned that bursty arrivals as well as time-varying service and arrival rates make estimation of parameters of the system very difficult and often inaccurate, and used this reason as a main motivation for designing robust scheduling policies. However, the theoretical results are provided for a time-invariant system with memoryless queues. In this section, we investigate the performance of our proposed algorithm in a time-varying system with bursty arrivals.

Consider the same DAG structure of previous simulations and the same server rates. We model the burstiness of demand as follows. At each time slot, a batch of B jobs arrive to the system with probability λ/B. To simulate a time-varying system, we consider two modes of network parameters. In the first mode, arrival rate is $\lambda = 1/5$, and task service rates are

$$\mu_1 = 1, \mu_2 = 4/3, \mu_3 = 2, \mu_4 = 1/2 \text{ and } \mu_5 = 2/3.$$

In the second mode, arrival rate is $\lambda = 1/6$, and task service rates are

$$\mu_1 = 1/2, \mu_2 = 2, \mu_3 = 1, \mu_4 = 2/5 \text{ and } \mu_5 = 1.$$

Note that the capacity region of mode 2 is $\lambda < 3/14$. We simulate a network that changes mode every T time slots.

Figure 4d illustrates the queue-length of the queue 4 versus time when the system has parameters $T = 1000$ and $B = 5$, and $\delta = 0.02$. As one expects, the bursty time-variant system suffers from larger delay. However, as the simulation shows the queue is still stable, and the gradient algorithm is able to track the changes in the network parameters.

2.6. Unstable Network with Generic Service Rates

In this subsection, we first propose a natural extension of the robust algorithm for the case that service rates are generic, using an approach that is similar to the design of policy when service rates satisfy Assumption 1.

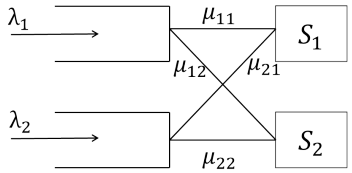

FIGURE 5: The X model

Define the allocation vector to be $p = [p_{kj}]$. Similar to before, the algorithm tries to minimize $\sum_{k=1}^{K}(\nu_k - \sum_{j=1}^{J}\mu_{kj}p_{kj})^2$ using gradient method. Then, a non-robust update of the allocation vector would be

$$p_{kj}^{n+1} = [p_{kj}^n + \beta^n \mu_{kj}(\nu_k - \sum_{j=1}^{J}\mu_{kj}p_{kj}^n)]_{\mathcal{C}}. \tag{16}$$

A robustified version of the update is

$$p_{kj}^{n+1} = [p_{kj}^n + \beta^n \mathbf{1}_{E_k^n} \sum_{(i',i)\in\mathcal{H}_k} \Delta Q_{(i',i)}^{n+1}]_{\mathcal{C}}. \tag{17}$$

We now demonstrate that the extension of our robust algorithm need not be throughput-optimal via a simple example. We consider a specific queueing network known in the literature as "X" model [5], shown in Figure 5 with generic service rates μ_{kj} that cannot be factorized into 2 factors μ_k and α_j. In our setting, the queueing network is equivalent to having 2 types of DAGs, each of them consisted of a single task with different service characteristics. There are 2 servers in the system. The network parameters are as follows.

$$\lambda_1 = \lambda_2 = 0.3,\ \mu_{11} = \mu_{22} = 1/8,\ \text{and}\ \mu_{12} = \mu_{21} = 3/8.$$

It is easy to check that the stability region is $\lambda_1 \leq 3/8$ and $\lambda_2 \leq 3/8$, which is achieved by server 1 always working on task 2 and server 2 always working on task 1. Let $p = [p_{11}, p_{12}, p_{21}, p_{22}]$ be the allocation vector for this example. Then, the update is

$$\begin{bmatrix} p_{1j}^{n+1} \\ p_{2j}^{n+1} \end{bmatrix} = \begin{bmatrix} p_{1j}^n + \beta^n \mathbf{1}_{E_1^n} \Delta Q_1^{n+1} \\ p_{2j}^n + \beta^n \mathbf{1}_{E_2^n} \Delta Q_2^{n+1} \end{bmatrix}_{\mathcal{C}}$$

for $j = 1, 2$. We remark that the update of allocation vector is identical for both servers. Thus, it is expected that the allocation vector converges to $p_{kj} = 1/2$ for all $k \in \{1, 2\}$ and $j \in \{1, 2\}$, due to symmetry. The simulation results confirm that

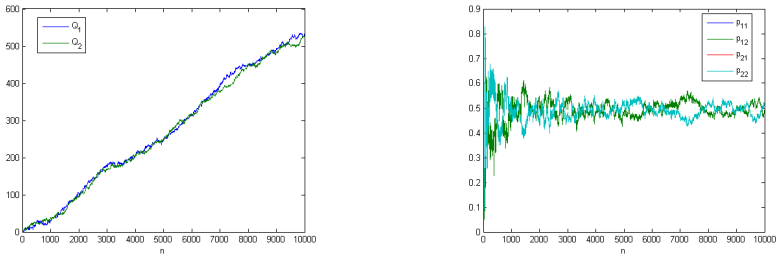

(a) The robust scheduling policy is unstable for the X model.

(b) Allocation vector p^n converges to $[1/2, 1/2, 1/2, 1/2]$.

FIGURE 6: This figure shows the simulation results for the X model.

our scheduling policy is not stable with these network parameters, since the allocation vector converges to $[1/2, 1/2, 1/2, 1/2]$. In this simulation, we set $p^0_{kj} = 0.1$ for all k and j. In general, the reason for the convergence to a sub-optimal allocation vector for generic μ_{kj} is that the skewed gradient projection (after dropping the term μ_{kj} from (16) to (17)) does not converge, even without noise.

In general, it would be interesting to find out whether there exists a robust scheduling policy that stabilizes the X model. Due to the underlying symmetry of the problem, it is reasonable to conjecture that no myopic queue-size policy[4] can be throughput-optimal in this example.

3. Flexible Queueing Network

In this section, we consider a different queueing network model, and show that our robust scheduling policy can also be applied to this network.

3.1. Network Model

We consider a flexible queueing network with K queues and J servers, and probabilistic routing. Servers are *flexible* in the sense that each server can serve a (non-empty) set of queues. Similarly, tasks in each queue are *flexible*, so that each queue can be served by a set of servers. Similar to the DAG scheduling model, for each j, let $\mathcal{T}_j$ be the set of queues that server j can serve, let $T_j = |\mathcal{T}_j|$. For each k, let $\mathcal{S}_k$ be the

[4]These are scheduling policies that are only a function of the current queue sizes of the network.

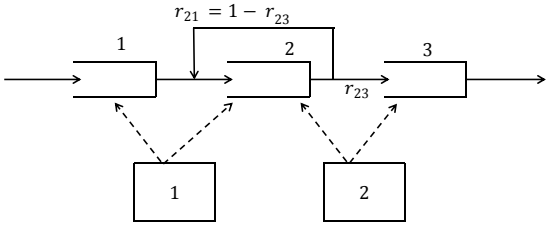

FIGURE 7: Flexible queueing network with 3 queues and 2 servers.

set of servers that can serve queue k, and let $S_k = |\mathcal{S}_k|$. Clearly, $\sum_{j=1}^{J} T_j = \sum_{k=1}^{K} S_k$, and we denote this sum by S. Without loss of generality, we assume that each server can serve at least one queue, and each queue can be served by at least one server.

We suppose that each queue has a dedicated exogenous arrival process (with rates being possibly zero). For each k, suppose that arrivals to queue k form an independent Bernoulli process with rate $\lambda_k \in [0, 1]$. Thus, in each time slot, there is exactly one arrival to queue k with probability λ_k, and no arrival with probability $1 - \lambda_k$. Let $A_k(t)$ be the cumulative number of exogenous arrivals to queue k up to time t. The routing structure of the network is described by the matrix $R = [r_{k'k}]_{1 \le k', k \le K}$, where $r_{k'k}$ denotes the probability that a task from queue k' joins queue k after service completion. The random routing is i.i.d. over all time slots. We assume that the network is *open*, i.e., all tasks eventually leave the system. This is characterized by the condition that $(I - R^T)$ is invertible, where I is the identity matrix, and R^T is the transpose of R.

Example 2. To clarify the network model, we consider a flexible queueing network shown in Figure 7. For concreteness, we can think of this system as a multi-tier application [23] with two flexible servers (the two boxes), and one type of application with three tiers in succession (the three queues). When a task is processed at queue 2, it will join queue 3 with probability r_{23} and it will join queue 1 with probability $r_{21} = 1 - r_{23}$ (that can be thought of as the failure probability in processing queue 2). This network is different from the classical open multiclass queueing networks, in that queue 2 can be served by 2 servers. In this network $\mathcal{T}_1 = \{1, 2\}$ and $\mathcal{T}_2 = \{2, 3\}$, $\mathcal{S}_1 = \{1\}$, $\mathcal{S}_2 = \{1, 2\}$, and $\mathcal{S}_3 = \{2\}$.

Similar to the DAG scheduling problem, we assume that several servers can work

simultaneously on the same task, so that their service capacities can be added. In each time slot, if a task in queue k is served exclusively by server j, then the task departs from queue k with probability $\mu_{kj} = \mu_j \alpha_j$, where μ_k is the service rate of queue k and α_j is the speed of server j.

The dynamics of the flexible queueing network can be described as follows. Let Q_k^n be the length of queue k at time n. Let $d_k^n \in \{0,1\}$ be the number of tasks that depart queue k at time n. Let $a_k^n \in \{0,1\}$ be the number of exogenous arrivals to queue k at time n. Finally, let $\mathbf{1}_{k\to k'}^n$ be the indicator that the task departing queue k at time n (if any) is destined to queue k'. Then the queue dynamics is

$$Q_k^{n+1} = Q_k^n + a_k^n + \sum_{k'=1}^{K} d_{k'}^n \mathbf{1}_{k'\to k}^n - d_k^n. \tag{18}$$

Note that $\mathbb{E}(a_k^n) = \lambda_k$ and $\mathbb{E}(\mathbf{1}_{k'\to k}^n) = r_{k'k}$.

Similar to the DAG scheduling problem, we define the allocation vector $p = [p_{kj}]$ (of server capacities), and p is called *feasible* if

$$\sum_{k\in\mathcal{T}_j} p_{kj} \le 1, \ \forall\ 1 \le j \le J. \tag{19}$$

For each j, p_{kj} can be interpreted as the probability that server j decides to work on queue k. Then, the head-of-the-line task in queue k is served with probability $\sum_j \mu_{kj} p_{kj}$. Similar to the DAG model, we scale the service rates so that $\sum_j \mu_{kj} p_{kj} < 1$ for any feasible p.

We now introduce the linear program (LP) that characterizes the capacity region of the flexible queueing network. Toward this end, for a given arrival rate vector λ, we first find the *nominal* traffic rates $\nu = [\nu_k]_{1\le k\le K} \in \mathbb{R}^K$, where ν_k is the long-run average total rate at which tasks arrive to queue k. For each k, $\nu_k = \lambda_k + \sum_{i=1}^K \nu_i r_{ik}$. Thus, we can solve ν in terms of R and λ:

$$\nu = (I - R^T)^{-1}\lambda. \tag{20}$$

Note that Eq. (20) is valid, since by our assumption that the network is open, $(I - R^T)$

is invertible. The LP is then defined as follows.

$$\text{Minimize} \quad \rho \tag{21}$$

$$\text{subject to} \quad \nu_k \leq \sum_{j \in \mathcal{S}_k} \mu_{kj} p_{kj}, \ \forall\ 1 \leq k \leq K, \tag{22}$$

$$\rho \geq \sum_{k \in \mathcal{T}_j} p_{kj}, \qquad \forall\ 1 \leq j \leq J, \tag{23}$$

$$p_{kj} = 0, \qquad \text{if } k \notin \mathcal{T}_j, \tag{24}$$

$$p_{kj} \geq 0. \tag{25}$$

Let the optimal value of the LP be ρ^*. Similar to Proposition 2.1 one can show that $\rho^* \leq 1$ is a necessary and sufficient condition of system stability. Thus, given μ_{kj} and R, the *capacity region* Λ of the network is the set of all $\lambda \in \mathbb{R}_+^K$, so that the corresponding optimal solution ρ^* to the LP satisfies $\rho^* \leq 1$.

3.2. Robust Scheduling Policy

In this section, we propose a robust scheduling policy that is provably throughput-optimal when the service rates can be written as $\mu_{kj} = \mu_k \alpha_j$. The policy is robust to arrival and task service rates, but not robust to routing probabilities of the network and servers' speed. The key idea is to use a stochastic gradient projection algorithm to update the service allocation vector p such that all the flows in the network are balanced. We first give the precise description of the algorithm, and state the main theorem. Then, we provide some explanations. We use similar notation as the one used in Section 2.

Since service rates can be factorized to a task-dependent rate and a server-dependent rate, only the sum $p_k \triangleq \sum_j \alpha_j p_{kj}$ affects the effective service rate for queue k. So, similar to the DAG scheduling problem, we call $p = [p_k] \in \mathbb{R}^K$ the service allocation vector. Our scheduling algorithm updates the allocation vector p^n in each time slot n in the following manner.

1. We initialize with an arbitrary feasible p^0.
2. Update the allocation vector p^n as follows.

$$p^{n+1} = [p^n + \beta^n \tilde{E}^n (I - R^T)^{-1} \Delta Q^n]_{\mathcal{C}_{\varepsilon_0}}, \tag{26}$$

where $\tilde{E}^n$ is a $K \times K$ diagonal matrix such that $\tilde{E}_{kk} = \mathbf{1}_{\{Q_k^n > 0\}}$ and $\mathcal{C}_{\varepsilon_0}$ is given in (14).

The main results of this section are the following two theorems.

Theorem 3.1. *Let $\lambda \in \Lambda$. The allocation vector p^n updated by (26) converges to $p^* = [p_k^*]$ almost surely, where $p_k^* = \frac{\nu_k}{\mu_k}$.*

The proof of Theorem 3.1 is almost identical to the proof of Theorem 2.1. We avoid repeating the details.

Theorem 3.2. *The flexible queueing network is rate stable under the robust scheduling algorithm specified by the update in (26), i.e.*

$$\lim_{n\to\infty} \frac{Q_k^n}{n} = 0, \ \forall k.$$

The proof of Theorem 3.2 is provided in Appendix 5.4.

The intuition for the update (26) is as follows. The algorithm tries to adaptively find the allocation vector p^* using a gradient projection method that solves (11). To robustify the algorithm to the knowledge of task service rates, we consider the "skewed" updates in (13). However, the major difference compared to the DAG scheduling problem is the way we find unbiased estimators of the terms $\nu_k - \mu_k p_k^n$. We use ΔQ^{n+1}, the changes in queue sizes, and routing matrix R, to estimate these terms. It is easy to show that the k^{th} entry of $(I - R^T)^{-1}\Delta Q^{n+1}$ is an unbiased estimator $\nu_k - \mu_k p_k^n$, if $Q_k^n > 0$. Define $M = \text{diag}\{\mu_k\}$, the $K \times K$ diagonal matrix with diagonal entries μ_k. Then,

$$\begin{aligned}
&\mathbb{E}(\tilde{E}^n(I - R^T)^{-1}\Delta Q^{n+1}|Q^n) && (27)\\
&\quad = \tilde{E}^n(I - R^T)^{-1}\mathbb{E}(\Delta Q^{n+1}|Q^n) && (28)\\
&\quad = \tilde{E}^n(I - R^T)^{-1}(\lambda + R^T M\tilde{E}^n p^n - M\tilde{E}^n p^n) && (29)\\
&\quad = \tilde{E}^n\nu - M\tilde{E}^n p^n && (30)\\
&\quad = \tilde{E}^n(\nu - Mp^n). && (31)
\end{aligned}$$

Note that matrix $\tilde{E}^n$ in update (26) ensures that the algorithm updates p_k^n only for queues k that are non-empty, since $\left[(I - R^T)^{-1}\Delta Q^{n+1}\right]_k$ is no longer an unbiased estimator of $\nu_k - \mu_k p_k^n$ when $Q_k^n = 0$.

4. Conclusion and Future Work

In this paper, we presented two processing networks that can well model different applications such as cloud computing, manufacturing lines, and healthcare systems. Our processing system is flexible in the sense that servers are capable of processing different types of tasks, while tasks can also be served by different servers. We proposed a scheduling and capacity allocation policy for these networks that is robust to service rates of the tasks and the arrival rates. The proposed scheduling algorithm is based on solving an optimization problem by stochastic gradient projection. The algorithm solves the problem of balancing all the flows in the network using only queue size information, and uses the allocation vector derived by the gradient algorithm at each time slot as its scheduling decision. We proved rate stability of the queueing networks corresponding to the models in the case that servers are cooperative and service rates can be factorized to a task-dependent rate and server-dependent rate.

There are many possible directions for future research. We summarize some of these directions as follows.

- It is important to find provably throughput optimal and robust scheduling policies while relaxing the assumption of cooperative servers. Some progress has been made in [25].
- A future direction is to investigate whether there exists a throughput-optimal policy which is only dependent on the queue-size information in the network in the current state or in the past, when service rates are generic, and cannot be necessary factorized to a task-dependent rate and server-dependent rate.
- In the case of flexible queueing networks, a future direction is to find a throughput-optimal policy that is robust to the knowledge of routing probabilities.

References

[1] Andradóttir, S., Ayhan, H. and Down, D. G. (2003). Dynamic server allocation for queueing networks with flexible servers. *Operations Research* **51,** 952–968.

[2] Atar, R., Mandelbaum, A. and Zviran, A. (2012). Control of fork-join

networks in heavy traffic. *50th Annual Allerton Conference on Communication, Control, and Computing* 823–830.

[3] Baccelli, F. and Makowski, A. (1985). Simple computable bounds for the fork-join queue. In *John Hopkins Conf. Information Science*.

[4] Baccelli, F., Makowski, A. M. and Shwartz, A. (1989). The fork-join queue and related systems with synchronization constraints: Stochastic ordering and computable bounds. *Advances in Applied Probability* 629–660.

[5] Baharian, G. and Tezcan, T. (2011). Stability analysis of parallel server systems under longest queue first. *Mathematical Methods of Operations Research* **74,** 257–279.

[6] Bambos, N. and Walrand, J. (1991). On stability and performance of parallel processing systems. *Journal of the ACM (JACM)* **38,** 429–452.

[7] Borkar, V. S. (2008). Stochastic approximation. *Cambridge Books*.

[8] Boyd, S. and Vandenberghe, L. (2004). *Convex optimization*. Cambridge university press.

[9] Chow, Y. S. and Teicher, H. (1997). *Probability Theory: Independence, Interchangeability, Martingales*. Springer Texts in Statistics.

[10] Dai, J. G. (1995). On positive harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. *The Annals of Applied Probability* 49–77.

[11] Dai, J. J. (1999). *Stability of fluid and stochastic processing networks*. University of Aarhus. Centre for Mathematical Physics and Stochastics (MaPhySto)[MPS].

[12] Dean, J. and Ghemawat, S. (2008). Mapreduce: simplified data processing on large clusters. *Communications of the ACM* **51,** 107–113.

[13] Dimakis, A. and Walrand, J. (2006). Sufficient conditions for stability of longest-queue-first scheduling: Second-order properties using fluid limits. *Advances in Applied probability* 505–521.

[14] Durrett, R. (2010). *Probability: theory and examples*. Cambridge university press.

[15] Harrison, J. M. (2000). Brownian models of open processing networks: Canonical representation of workload. *Annals of Applied Probability* 75–103.

[16] Harrison, J. M. and Nguyen, V. (1993). Brownian models of multiclass queueing networks: Current status and open problems. *Queueing Systems* **13,** 5–40.

[17] Jiang, L. and Walrand, J. (2010). A distributed csma algorithm for throughput and utility maximization in wireless networks. *IEEE/ACM Transactions on Networking (TON)* **18,** 960–972.

[18] Kandula, S., Sengupta, S., Greenberg, A., Patel, P. and Chaiken, R. (2009). The nature of data center traffic: measurements & analysis. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*. ACM. pp. 202–208.

[19] Konstantopoulos, P. and Walrand, J. (1989). Stationary and stability of fork-join networks. *Journal of Applied Probability* 604–614.

[20] Mandelbaum, A. and Stolyar, A. L. (2004). Scheduling flexible servers with convex delay costs: Heavy-traffic optimality of the generalized cμ-rule. *Operations Research* **52,** 836–855.

[21] Nguyen, V. (1993). Processing networks with parallel and sequential tasks: Heavy traffic analysis and brownian limits. *The Annals of Applied Probability* 28–55.

[22] Nguyen, V. (1994). The trouble with diversity: Fork-join networks with heterogeneous customer population. *The Annals of Applied Probability* 1–25.

[23] Padala, P., Hou, K.-Y., Shin, K. G., Zhu, X., Uysal, M., Wang, Z., Singhal, S. and Merchant, A. (2009). Automated control of multiple virtualized resources. In *Proceedings of the 4th ACM European conference on Computer systems*. ACM. pp. 13–26.

[24] Pedarsani, R., Walrand, J. and Zhong, Y. (2014). Robust scheduling in a flexible fork-join network. In *IEEE Conference on Decision and Control (CDC)*. IEEE. pp. 3669–3676.

[25] Pedarsani, R., Walrand, J. and Zhong, Y. (2014). Scheduling tasks with precedence constraints on multiple servers. *52nd Annual Allerton Conference on Communication, Control, and Computing* 1196–1203.

[26] Stolyar, A. L. and Yudovina, E. (2012). Tightness of invariant distributions of a large-scale flexible service system under a priority discipline. *Stochastic Systems* **2,** 381–408.

5. Appendix

5.1. Proof of Proposition 2.1

Consider the fluid scaling of the queueing network, $X^{rt}_{(k',k)} = \frac{Q_{(k',k)}(\lfloor rt \rfloor)}{r}$ (see [10] for more discussion on the stability of fluid models), and let $X^t_{(k',k)}$ be the corresponding fluid limit. The fluid model dynamics is as follows. If k is a root node, then

$$X^t_{(0,k)} = X^0_{(0,k)} + A^t_{m(k)} - D^t_k,$$

where $A^t_{m(k)}$ is the total number of jobs of type m (scaled to the fluid level) that have arrived to the system until time t. If k is not a root node, then,

$$X^t_{(k',k)} = X^0_{(k',k)} + D^t_{k'} - D^t_k,$$

where D^t_k is the total number of tasks (scaled to the fluid level) of type k processed up to time t. Suppose that $\rho^* > 1$. We show that if $X^0_{(k',k)} = 0$ for all (k',k), there exists t_0 and (k',k) such that $X^{t_0}_{(k',k)} \geq \epsilon(t_0) > 0$, which implies that the system is weakly unstable [11]. In contrary suppose that there exists a scheduling policy that under that policy for all $t \geq 0$ and all (k',k), $X^t_{(k',k)} = 0$. Pick a regular point[5] t_1. Then, for all (k',k), $\dot{X}^{t_1}_{(k',k)} = 0$. Since $\dot{A}^{t_1}_{m(k)} = \lambda_m = \nu_k$, this implies that $\dot{D}^{t_1}_k = \nu_k$ for all the root nodes k. Now considering queues (k',k) such that nodes k' are roots, one gets

$$\dot{D}^{t_1}_k = \dot{D}^{t_1}_{k'} = \nu_{k'} = \nu_k.$$

Similarly, one can inductively show that for all k, $\dot{D}^{t_1}_k = \nu_k$. On the other hand, at a regular point t_1, $\dot{D}^{t_1}_k$ is exactly the total service capacity allocated to task k at t_1. This implies that there exists p_{kj} at time t_1 such that $\nu_k = \sum_{j \in \mathcal{S}_k} \mu_{kj} p_{kj}$ for all k and the allocation vector $[p_{kj}]$ is feasible, i.e. $\sum_{k \in \mathcal{T}_j} p_{kj} \leq 1$. This contradicts $\rho^* > 1$.

Now suppose that $\rho^* \leq 1$, and $p^* = [p^*_{kj}]$ is an allocation vector that solves the LP. To prove sufficiency of the condition, consider a generalized head-of-the-line processor sharing policy that server j works on task k with capacity p^*_{kj}. Then the cumulative service allocated to task k up to time t is $S^t_k = \sum_{j \in \mathcal{S}_k} \mu_{kj} p^*_{kj} t \geq \nu_k t$. We show that $X^t_{(k',k)} = 0$ for all t and all (k',k), if $X^0_{(k',k)} = 0$ for all (k',k). First consider queue $(0,k)$ corresponding to a root node. Suppose that $X^{t_1}_{(0,k)} \geq \epsilon > 0$ for some

[5] We define a point t to be regular if $X^t_{(k',k)}$ is differentiable at t for all (k',k).

positive t_1 and ϵ. By continuity of the fluid limit, there exists $t_0 \in (0, t_1)$ such that $X^{t_0}_{(0,k)} = \epsilon/2$ and $X^t_{(0,k)} > 0$ for all $t \in [t_0, t_1]$. Then, $\dot{X}^t_{(0,k)} = \nu_k - \sum_{j \in \mathcal{S}_k} \mu_{kj} p^*_{kj} \leq 0$ for $t \in [t_0, t_1]$, which is a contradiction. Now we show that $X^t_{(k',k)} = 0$ for all t if k' is a root node and k is a child of k'. Note that $X^t_{(0,k')} = 0$; thus, $\dot{D}^t_{k'} = \nu_{k'} = \nu_k$. Then, $\dot{X}_{(k',k)} = \nu_k - \sum_{j \in \mathcal{S}_k} \mu_{kj} p^*_{kj} \leq 0$. This proves that $X^t_{(k',k)} = 0$ for all t. One can then inductively complete this proof for all queues (k', k).

5.2. Proof of Theorem 2.1

Recall the following notation which will be widely used in the proofs.

$$\mathbf{1}_{\{Q^n_{(k',k)} > 0,\ \forall k' \in \mathcal{P}_k\}} = \mathbf{1}_{E^n_k}.$$

We introduce another notation which is

$$\mathbf{1}_{E^n} = \prod_{k=1}^{K} \mathbf{1}_{E^n_k}.$$

Note that event E^n denotes the event that all the queues are non-empty at time n.

Lemma 5.1. *There exist constants ℓ and $\delta_0 > 0$, which are independent of n, such that given any history $\mathcal{F}^n$ up to time n, $\mathbb{P}(\mathbf{1}_{E^{n+\ell}} = 1 | \mathcal{F}^n) \geq \delta_0 > 0$.*

Proof. We work with each of the DAGs $\mathcal{G}_m$ separately, and construct events so that all the queues corresponding to $\mathcal{G}_m$ have positive lengths after some time ℓ. We can do this since $\mu_k p^n_k$ will always be no smaller than $\mu_k \varepsilon_0$ and strictly smaller than 1, so there is positive probability of serving or not serving a task.

Let $\widetilde{E}^n_k$ be the event that task k is served at time n, $\bar{E}^n_k$ be the event that task k is not served at time n, and $\hat{E}^n_m$ be the event that job type m arrives at time n. Consider a particular DAG $\mathcal{G}_m$. Recall that L_k is the length of the longest path from the root nodes of the DAG to node k. Let $\ell_m = \max_{k \in \mathcal{V}_m} L_k + 1$. We construct the event $E(\ell_m)$ that happens with a strictly positive probability, and assures that all the queues at time $n + \ell_m$ are non-empty. Toward this end, let $E(\ell_m) = \cap_{n'=0}^{\ell_m - 1} C^{n'}$, where event $C^{n'}$ is

$$C^{n'} = \hat{E}^{n'}_m \cap_{\{k: L_k \leq n'-1\}} \widetilde{E}^{n'}_k \cap_{\{k: L_k > n'-1\}} \bar{E}^{n'}_k.$$

In words, $C^{n'}$ is the event that at time n', there is a job arrival of type m, services of tasks of class k for k with $L_k \leq n' - 1$, and no service of tasks of class k for k with

$L_k > n' - 1$. Now, by construction all the queues are non-empty at time $n + \ell$ with a positive probability. To illustrate how we construct this event, consider the example of Figure 2a and the corresponding queueing network in Figure 2b. Then, C^0 is the event that there is an arrival to the system, and no service in the network. C^1 is the event that there is a new job arriving, task 1 is served, and tasks 2, 3, and 4 are not served. Note that there is certainly at least one available task 1 to serve due to C^0. Up to now, certainly queues $(0, 1)$, $(1, 2)$, and $(1, 3)$ are non-empty. C^2 is the event of having a new arrival, service to tasks 1, 2, and 3, and no service to task 4. This construction assures that after 3 time slots, all the queues are non-empty.

Now for the whole network it is sufficient to take $\ell = \max_m \ell_m$. Construct the events $E(\ell_m)$ for each DAG independently, and freeze the DAG $\mathcal{G}_m$ (no service and no arrivals) from time $n + \ell_m$ to $n + \ell - 1$. This construction makes sure that all the queues in the network are non-empty at time $n + \ell$ given any history $\mathcal{F}^n$ with some positive probability δ_0.

Lemma 5.2. *The following inequality holds.*

$$\liminf_{n\to\infty} \frac{1}{n} \sum_{n'=1}^{n} \mathbf{1}_{E^{n'}} \geq \frac{\delta_0}{\ell} > 0, a.s.$$

Proof. Take a subsequence $\mathbf{1}_{E^{n'\ell}},\ n' \geq 1$. Define a sequence $\left(Y^{n'}\right)_{n'\geq 1}$ by

$$Y^{n'} = \mathbf{1}_{E^{n'\ell}} - \mathbb{P}(\mathbf{1}_{E^{n'\ell}} = 1 \mid \mathcal{F}^{(n'-1)\ell}).$$

Then, it is easy to see that $\mathbb{E}\left[Y^{n'} \mid \mathcal{F}^{(n'-1)\ell}\right] = 0$. Thus, $\left(Y^{n'}\right)$ is an $\mathcal{F}^{n'\ell}$-adapted zero-mean martingale. Furthermore, we have $|Y^{n'}| \leq 2$ a.s. for each n'. By applying the martingale law of large numbers (see e.g., Corollary 2 in Section 11.2 of [9]), we have $\lim_{m\to\infty} \frac{1}{m} \sum_{n'=1}^{m} Y^{n'} = 0$ a.s. By Lemma 5.1, this immediately implies that $\liminf_{m\to\infty} \frac{1}{m} \sum_{n'=1}^{m} \mathbf{1}_{E^{n'\ell}} \geq \delta_0$ a.s. Therefore, with probability 1,

$$\liminf_{n\to\infty} \frac{1}{n} \sum_{n'=1}^{n} \mathbf{1}_{E^{n'}} \geq \liminf_{n\to\infty} \frac{1}{n} \sum_{n'=1}^{\lfloor n/\ell \rfloor} \mathbf{1}_{E^{n'\ell}} \geq \frac{\delta_0}{\ell}.$$

This completes the proof of Lemma 5.2.

Lemma 5.3. *The following equality holds.*

$$\lim_{n\to\infty} \sum_{n'=1}^{n} \beta^{n'} \mathbf{1}_{E^{n'}} = \infty,\ a.s.$$

Proof. From now on we work with the probability-1 event defined in Lemma 5.2. Consider a sample path in this probability-1 event, and let $x_{n'} = \beta^{n'} \mathbf{1}_{E^{n'}}$. First note that $x_{n'} \geq 0$. Thus, by the monotone convergence theorem, the series either converges or goes to infinity. Suppose that $\lim_{n\to\infty} \sum_{n'=1}^{n} x_{n'} = s$ for some finite s. Define the sequence $b_{n'} = \frac{1}{\beta^{n'}}$. Then, by Kronecker's lemma [14], we have $\lim_{n\to\infty} \frac{1}{b_n} \sum_{n'=1}^{n} b_{n'} x_{n'} = 0$. This shows that $\lim_{n\to\infty} \frac{1}{b_n} \sum_{n'=1}^{n} \mathbf{1}_{E^{n'}} = 0$, which results in a contradiction, since $\lim_{n\to\infty} \frac{1}{n\beta^n}$ is finite, and hence by Lemma 5.2,

$$\liminf_{n\to\infty} \frac{1}{b_n} \sum_{n'=1}^{n} \mathbf{1}_{E^{n'}} > 0.$$

Now we are ready to prove Theorem 2.1. Consider the probability-1 event in Lemma 5.3. Let $d_n = \|p^n - p^*\|^2$ and fix $\epsilon > 0$. We prove that there exists a $n_0(\epsilon)$ such that for all $n \geq n_0(\epsilon)$, d_n has the following properties.

(i) If $d_n < \epsilon$, then $d_{n+1} < 3\epsilon$.

(ii) If $d_n \geq \epsilon$ then, $d_{n+1} \leq d_n - \gamma^n$ where $\sum_{n=1}^{\infty} \gamma^n = \infty$ and $\gamma^n \to 0$.

Then property (ii) shows that for some large enough $n_1 = n_1(\epsilon) > n_0(\epsilon)$, $d_{n_1} < \epsilon$, and properties (i) and (ii) show that $d_n < 3\epsilon$ for $n \geq n_1(\epsilon)$. This is true for all $\epsilon > 0$, so d_n converges to 0 almost surely.

First we show property (i). Let $U^n = [U_k^n] \in \mathbb{R}^K$ be the vector of updates such that

$$U_k^{n+1} = \mathbf{1}_{E_k^n} \sum_{(i',i)\in\mathcal{H}_k} \Delta Q_{(i',i)}^{n+1}.$$

Note that $\|U^n\|^2$ is bounded by some constant $C_1 > 0$, since the queues length changes at each time slot are bounded by 1. On the other hand, $\beta^n \to 0$. Thus, one can take $n_1(\epsilon)$ large enough such that for all $n \geq n_1(\epsilon)$, $\beta^n \leq \sqrt{\frac{\epsilon}{2C_1}}$. Then, for $n \geq n_1(\epsilon)$ if $d_n < \epsilon$,

$$\begin{aligned} d_{n+1} &= \|p^{n+1} - p^*\|^2 && (32)\\ &= \|[p^n + \beta^n U^{n+1}]_{\mathcal{C}_\varepsilon} - p^*\|^2 && (33)\\ &\leq \|p^n + \beta^n U^{n+1} - p^*\|^2 && (34)\\ &\leq 2d_n + 2(\beta^n)^2 \|U^{n+1}\|^2 < 3\epsilon, && (35) \end{aligned}$$

where (34) is due to the fact that projection to the convex set is non-expansive, and (35) is by Cauchy-Schwarz inequality.

To show property (ii), we make essential use of the fact that the cumulative stochastic noise is a martingale. Let

$$Z_k^{n+1} = \mathbf{1}_{E_k^n}(U_k^n - \nu_k + \mu_k p_k^n).$$

Then, by (9),

$$\mathbb{E}\left[Z_k^{n+1}|\mathcal{F}^n\right] = 0, \ \forall k \tag{36}$$

which shows that Z^n is a martingale difference sequence. Now observe that

$$d_{n+1} = \sum_{k=1}^{K}(p_k^{n+1} - p_k^*)^2 \tag{37}$$

$$\leq \sum_{k=1}^{K}(p_k^n + \beta^n U_k^{n+1} - p_k^*)^2 \tag{38}$$

$$= d_n + (\beta^n)^2\|U^{n+1}\|^2 + 2\beta^n \sum_{k=1}^{K}(p_k^n - p_k^*)(\nu_k - \mu_k p_k^n + Z_k^{n+1})\mathbf{1}_{E_k^n} \tag{39}$$

$$= d_n + (\beta^n)^2\|U^{n+1}\|^2 + 2\beta^n \sum_{k=1}^{K}(p_k^n - p_k^*)(\mu_k(p_k^* - p_k^n) + Z_k^{n+1})\mathbf{1}_{E_k^n} \tag{40}$$

$$\leq d_n + (\beta^n)^2 C_1 - \sum_{k=1}^{K} 2\mu_k\beta^n \mathbf{1}_{E^n}(p_k^n - p_k^*)^2 + \sum_{k=1}^{K} 2\beta^n Z_k^{n+1}(p_k^n - p_k^*)\mathbf{1}_{E_k^n}, \tag{41}$$

where (38) is due to non-expansiveness of projection, and (41) is due the facts that $E^n \subset E_k^n$ and $\|U^n\|^2 \leq C_1$. Let $\mu^* = \min_k \mu_k$. Since $\sum_{k=1}^{K}(p_k^n - p_k^*)^2 > \epsilon$, the following choice of γ^n satisfies $d_{n+1} \leq d_n - \gamma^n$:

$$\gamma^n = -(\beta^n)^2 C_1 + \beta^n \mathbf{1}_{E^n} 2\mu^*\epsilon - \sum_{k=1}^{K} 2\beta^n Z_k^{n+1}(p_k^n - p_k^*)\mathbf{1}_{E_k^n}.$$

As $\beta^n \to 0$ as $n \to \infty$, it is easy to see that $\gamma^n \to 0$ almost surely. Thus, to complete the proof of Theorem 2.1, one needs to show that $\sum_{n=1}^{\infty}\gamma_n = \infty$ almost surely. Toward this end, note that $\sum_n(\beta^n)^2$ is finite which makes $-\sum_n(\beta^n)^2 C_1$ bounded. By (36), and the facts that $\sum_n(\beta^n)^2 < \infty$ and $\|p^n - p^*\|$ is bounded for all n, we get that

$$V^n = \sum_{n'=1}^{n}\sum_{k=1}^{K} 2\beta^{n'} Z_k^{n'+1}(p_k^{n'} - p_k^*)\mathbf{1}_{E_k^{n'}}$$

is an L_2-bounded martingale and by the martingale convergence theorem converges to some bounded random variable almost surely [14]. Finally,

$$2\epsilon\mu^* \sum_{n=1}^{\infty}\beta^n \mathbf{1}_{E^n} = \infty, a.s.$$

by Lemma 5.3. This completes the proof of Theorem 2.1.

5.3. Proof of Theorem 2.2

In this subsection, we provide the proof of Theorem 2.2. The key idea to prove the rate stability of queues is to first show that the servers allocate enough cumulative capacity to all the tasks in the network. This is formalized in Lemma 5.4. Second, in Lemma 5.5, we show that each queue (k', k) cannot go unstable if task k receives enough service allocation over time, and the traffic rate coming to these queues is nominal. Finally, we use these two conditions to show rate stability of all the queues in the network by mathematical induction.

To prove the theorem, we first introduce some notation. Let D_k^n denote the cumulative number of processed tasks of type k at time n. Recall that d_k^n is the number of processed tasks of type k at time n. Therefore, $D_k^n = \sum_{n'=1}^{n} d_k^{n'}$. Let $A_m^n = \sum_{n'=1}^{n} a_m^{n'},\ 1 \leq m \leq M$ be the cumulative number of jobs of type m that have arrived up to time n. Then the queue-length dynamic of queue (k', k) can be written as follows. If $k' \neq 0$, then $Q_{(k',k)}^n = Q_{(k',k)}^0 + D_{k'}^n - D_k^n$. If $k' = 0$, then $Q_{(k',k)}^n = Q_{(k',k)}^0 + A_{m(k)}^n - D_k^n$.

At time n, the probability that one task is served and departed from queue (k', k) is $\mu_k p_k^n$, if all of the queues (i, k) are non-empty for all i. We define s_k^n to be a random variable denoting the virtual service that queues (k', k) have received at time n, whether there has been an available task k to be processed or not. s_k^n is a Bernoulli random variable with parameter $\mu_k p_k^n$. Then, the cumulative service that queues (k', k) receive up to time n is $S_k^n = \sum_{n'=1}^{n} s_k^{n'}$ for all k'. Note that the cumulative service is different from the cumulative departure. Indeed, $d_k^n = s_k^n \mathbf{1}_{E_k^n}$.

From now on, in the proof of Theorem 2.2, we consider the probability-1 event that p^n converges to p^* stated in Theorem 2.1.

Lemma 5.4. *The following equality holds:*

$$\lim_{n\to\infty} \frac{S_k^n}{n} = \nu_k,\ a.s.,\ \forall k. \tag{42}$$

Proof. By Theorem 2.1, the sequence p_k^n converges to $\frac{\nu_k}{\mu_k}$ almost surely. Therefore, for all the sample paths in the probability-1 event, and for all $\epsilon_1 > 0$, there exists $n_0(\epsilon_1)$ such that $\|\mu_k p_k^n - \nu_k\| \leq \epsilon_1$, for all $n > n_0(\epsilon_1)$.

Let $\tilde{s}_k^n$ be i.i.d Bernoulli process of parameter $\nu_k - \epsilon_1$. We couple the processes s_k^n and $\tilde{s}_k^n$ as follows. If $s_k^n = 0$, then $\tilde{s}_k^n = 0$. If $s_k^n = 1$, then $\tilde{s}_k^n = 1$ with probability $\frac{\nu_k - \epsilon_1}{\mu_k p_k^n}$, and $\tilde{s}_k^n = 0$ with probability $1 - \frac{\nu_k - \epsilon_1}{\mu_k p_k^n}$. Note that $\tilde{s}_k^n$ is still marginally i.i.d Bernoulli process of parameter $\nu_k - \epsilon_1$. Then,

$$\begin{aligned}
\liminf_{n\to\infty} \frac{S_k^n}{n} &\geq \liminf_{n\to\infty} \frac{\sum_{n'=n_0(\epsilon_1)+1}^{n} s_k^{n'}}{n} && (43)\\
&\geq \liminf_{n\to\infty} \frac{\sum_{n'=n_0(\epsilon_1)+1}^{n} \tilde{s}_k^{n'}}{n} && (44)\\
&= \nu_k - \epsilon_1 \ a.s., && (45)
\end{aligned}$$

where (44) is by construction of the coupled sequences, and (45) is by the strong law of large numbers.

Let $\bar{s}_k^n$ be i.i.d Bernoulli process of parameter $\nu_k + \epsilon_1$. We couple the processes s_k^n and $\bar{s}_k^n$ as follows. If $s_k^n = 1$, then $\tilde{s}_k^n = 1$. If $s_k^n = 0$, then $\tilde{s}_k^n = 0$ with probability $\frac{1-(\nu_k+\epsilon_1)}{1-\mu_k p_k^n}$, and $\tilde{s}_k^n = 1$ with probability $1 - \frac{1-(\nu_k+\epsilon_1)}{1-\mu_k p_k^n}$. Note that $\bar{s}_k^n$ is still marginally i.i.d Bernoulli process of parameter $\nu_k + \epsilon_1$. Then,

$$\begin{aligned}
\limsup_{n\to\infty} \frac{S_k^n}{n} &\leq \limsup_{n\to\infty} \frac{n_0(\epsilon_1) + \sum_{n'=n_0(\epsilon_1)+1}^{n} s_k^{n'}}{n} && (46)\\
&\leq \limsup_{n\to\infty} \frac{n_0(\epsilon_1) + \sum_{n'=n_0(\epsilon_1)+1}^{n} \bar{s}_k^{n'}}{n} && (47)\\
&= \nu_k + \epsilon_1 \ a.s., && (48)
\end{aligned}$$

where (47) is by construction of the coupled sequences, and (48) is by the strong law of large numbers. Letting $\epsilon_1 \to 0$, we have

$$\liminf_{n\to\infty} \frac{S_k^n}{n} = \limsup_{n\to\infty} \frac{S_k^n}{n} = \nu_k, \ a.s.$$

Lemma 5.5. *Consider a fixed k and all queues (k', k) with $k' \in \mathcal{P}_k$. Suppose that*

$$\lim_{n\to\infty} \frac{D_{k'}^n}{n} = \nu_k, \ \forall k' \in \mathcal{P}_k, \ a.s. \tag{49}$$

if $\mathcal{P}_k \neq \emptyset$, and

$$\lim_{n\to\infty} \frac{A_{m(k)}^n}{n} = \nu_k, \ a.s. \tag{50}$$

if k is a root node. Then,

$$\lim_{n\to\infty} \frac{Q^n_{(k',k)}}{n} = 0, a.s.$$

Proof. Before getting to the details of the proof, note that if k is a root node, then we readily know that

$$\lim_{n\to\infty} \frac{A^n_{m(k)}}{n} = \lambda_m = \nu_k, a.s.$$

Thus, the lemma states that queues $(0,k)$ are rate stable.

We prove the lemma for the general case that node k is not a root node. Similar proof holds for the case of root nodes. First, we show that for all pair of queues (i,k) and (i',k) such that $i,i' \in \mathcal{P}_k$, we have

$$\lim_{n\to\infty} \frac{Q^n_{(i,k)} - Q^n_{(i',k)}}{n} = 0, a.s. \tag{51}$$

Note that

$$\frac{Q^n_{(i,k)} - Q^n_{(i',k)}}{n} = \frac{D^n_i - D^n_k - (D^n_{i'} - D^n_k)}{n} = \frac{D^n_{i'} - D^n_i}{n}.$$

Then, by (49),

$$\lim_{n\to\infty} \frac{D^n_i - D^n_{i'}}{n} = \nu_k - \nu_k = 0, a.s.$$

Second, we show that

$$\liminf_{n\to\infty} \frac{Q^n_{(k',k)}}{n} = 0, a.s.$$

In contrary suppose that in one realization in the probability-1 event defined by (49) and Lemma 5.4,

$$\liminf_{n\to\infty} \frac{Q^n_{(k',k)}}{n} > 2\epsilon_2, \tag{52}$$

for some $\epsilon_2 > 0$. This implies that in that realization,

$$\liminf_{n\to\infty} \frac{Q^0_{(k',k)} + D^n_{k'} - D^n_k}{n} > 2\epsilon_2.$$

By (49), the probability that $\lim_{n\to\infty} \frac{Q^0_{(k',k)} + D^n_{k'}}{n} = \nu_k$ is 1. Thus, in that realization

$$\limsup_{n\to\infty} \frac{D^n_k}{n} < \nu_k - 2\epsilon_2. \tag{53}$$

On the other hand, (52) shows that there exists $n_0(\epsilon_2)$ such that for all $n \geq n_0(\epsilon_2)$,

$$Q^n_{(k',k)} \geq 2n\epsilon_2. \tag{54}$$

Furthermore, (51) shows that there exists $n_1(\epsilon_2)$ such that for all $i \in \mathcal{P}_k$ and for all $n \geq n_1(\epsilon)$,

$$|Q^n_{(k',k)} - Q^n_{(i,k)}| < n\epsilon_2. \tag{55}$$

Let $n_2(\epsilon_2) = \max(n_0(\epsilon_2), n_1(\epsilon_2))$. (54) and (55) imply that for all $n \geq n_2(\epsilon_2)$, $Q^n_{(i,k)} \geq n\epsilon_2$. Now taking $n_3(\epsilon_2) = \max(n_2(\epsilon_2), 1/\epsilon_2)$, we have that all the queues $(i,k),\ i \in \mathcal{P}_k$ are non-empty for $n \geq n_3(\epsilon_2)$. Thus, $s^n_k = d^n_k$ for all $n \geq n_3(\epsilon_2)$. Therefore,

$$\limsup_{n\to\infty} \frac{S^n_k}{n} \leq \limsup_{n\to\infty} \frac{n_3(\epsilon_2) + D^n_k}{n} = \limsup_{n\to\infty} \frac{D^n_k}{n}.$$

Thus, by Lemma 5.4, $\nu_k \leq \limsup_{n\to\infty} \frac{D^n_k}{n}$ which contradicts (53). Since this holds for any $\epsilon_2 > 0$, we conclude that

$$\liminf_{n\to\infty} \frac{Q^n_{(k',k)}}{n} = 0,\ a.s., \tag{56}$$

for all $k' \in \mathcal{P}_k$.

Third, we show that

$$\limsup_{n\to\infty} \frac{Q^n_{(k',k)}}{n} = 0, a.s.$$

In contrary suppose that in one realization in the probability-1 event defined by (49), (51), and Lemma 5.4,

$$\limsup_{n\to\infty} \frac{Q^n_{(k',k)}}{n} > 4\epsilon_3, \tag{57}$$

for some $\epsilon_3 > 0$. This implies that in that realization $Q^n_{(k',k)} > 4n\epsilon_3$ happens infinitely often. Moreover, by (56), $Q^n_{(k',k)} < 2n\epsilon_3$ happens also infinitely often in that realization. On the other hand, by (51), there exists some $n_0(\epsilon_3)$ such that for all $n \geq n_0(\epsilon_3)$ and all $i \in \mathcal{P}_k$,

$$|Q^n_{(i,k)} - Q^n_{(k',k)}| < n\epsilon_3. \tag{58}$$

Take $N_1 = \max(n_0(\epsilon_3), \frac{2}{\epsilon_3})$. Then, there exists $N_1 \leq n_1(\epsilon_3) < n_2(\epsilon_3)$ such that $Q^{n_1}_{(k',k)} \leq 2n_1\epsilon_3$ and $Q^{n_2}_{(k',k)} \geq 4n_2\epsilon_3$. In words, $n_1 + 1$ is the first time after N_1 that $\frac{Q^n_{(k',k)}}{n}$ crosses $2\epsilon_3$ without going below $2\epsilon_3$ before exceeding $4\epsilon_3$. Then, since the queue-length changes by at most 1 each time slot, queue (k',k) is non-empty for all $n,\ n_1 \leq n \leq n_2$. Furthermore, for all n, $n_1 \leq n \leq n_2$ and for all $i \in \mathcal{P}_k$, by (58),

$$Q^n_{(i,k)} > Q^n_{(k',k)} - n\epsilon_3 \geq 2n\epsilon_3 - 1 - n\epsilon_3 \geq n_1\epsilon_3 - 1 \geq 1.$$

Thus, all the queues (i,k) are also non-empty for all n in the interval $n_1 \leq n \leq n_2$. Consequently, for all $n,\ n_1 \leq n \leq n_2,\ s_k^n = d_k^n$. Now define a process

$$B^n_{(k',k)} = D^n_{k'} - S^n_k.$$

Note that by (49) and Lemma 5.4, in the realization of probability-1 event that we consider,

$$\lim_{n\to\infty} \frac{B^n_{(k',k)}}{n} = \nu_k - \nu_k = 0. \tag{59}$$

We bound $B^{n_2}_{(k',k)}$ as follows.

$$\begin{aligned}
B^{n_2}_{(k',k)} &= B^{n_1}_{(k',k)} + [B^{n_2}_{(k',k)} - B^{n_1}_{(k',k)}] \\
&= B^{n_1}_{(k',k)} + [D^{n_2}_{k'} - D^{n_1}_{k'} - (S^{n_2}_k - S^{n_1}_k)] \\
&= B^{n_1}_{(k',k)} + [D^{n_2}_{k'} - D^{n_1}_{k'} - (D^{n_2}_k - D^{n_1}_k)] \\
&= B^{n_1}_{(k',k)} + [Q^{n_2}_k - Q^{n_1}_k] \\
&\geq B^{n_1}_{(k',k)} + 4\epsilon_3 n_2 - 2\epsilon_3 n_1.
\end{aligned}$$

Dividing both sides of the inequality by n_2 and subtracting $B^{n_1}_{(k',k)}/n_1$, one gets

$$\frac{B^{n_2}_{(k',k)}}{n_2} - \frac{B^{n_1}_{(k',k)}}{n_1} \geq \frac{B^{n_1}_{(k',k)}}{n_1}(\frac{n_1}{n_2} - 1) + 4\epsilon_3 - 2\epsilon_3 \frac{n_1}{n_2}.$$

By (59), one can choose a large enough N_2 such that for all $n \geq N_2$, $|\frac{B^n_{(k',k)}}{n}| \leq \frac{2\epsilon_3}{3}$. Then, N_1 can be chosen as

$$N_1 = \max(n_0(\epsilon_3), \frac{2}{\epsilon_3}, N_2),$$

and one chooses n_1 and n_2 accordingly as before. Then, since $n_1, n_2 \geq N_1$, one can write

$$|\frac{B^n_{(k',k)}}{n} - \frac{B^n_{(k',k)}}{n}| \leq \frac{4\epsilon_3}{3}. \tag{60}$$

However,

$$\frac{B^{n_2}_{(k',k)}}{n_2} - \frac{B^{n_1}_{(k',k)}}{n_1} \geq 2\epsilon_3(\frac{n_1}{n_2} - 1) + 4\epsilon_3 - 2\epsilon_3 \frac{n_1}{n_2} = 2\epsilon_3,$$

which contradicts (60). Thus,

$$\limsup_{n\to\infty} \frac{Q^n_{(k',k)}}{n} = 0, a.s.$$

The result holds for arbitrary $k' \in \mathcal{P}_k$. This completes the proof of Lemma 5.5.

Now we are ready to prove Theorem 2.2. We complete the proof of Theorem 2.2 by induction. Recall that L_k is the length of the longest path from the root of the DAG $\mathcal{G}_{m(k)}$ to node k. If k is a root, $L_k = 0$. The formal induction goes as follows.

- **Basis:** All the queues corresponding to root nodes, i.e. all (k', k) for which $L_k = 0$ are rate stable.
- **Inductive Step:** If all the queues (k', k) for which $L_k \leq L - 1$ are rate stable, then all the queues (k', k) for which $L_k = L$ are also rate stable.

The basis is true by Lemma 5.5. The inductive step is also easy to show using Lemma 5.5. For a particular queue $(k', k) = (i_L, i_{L+1})$, suppose that $L_k = L$. Pick a path of edges

$$(i_1, i_2), (i_2, i_3), \ldots, (i_L, i_{L+1}),$$

from queue $(0, i_1)$ to (k', k). By assumption of induction, all the queues (i_l, i_{l+1}) are rate stable for $l \leq L - 1$.

$$\begin{aligned} A^n_{m(k)} - D^n_{i_L} &= A^n_{m(k)} - D^n_{i_1} + \sum_{l=1}^{L-1} (D^n_{i_l} - D^n_{i_{l+1}}) \\ &= Q^n_{(0,i_1)} - Q^0_{(0,i_1)} + \sum_{l=1}^{L-1} (Q^n_{(i_l,i_{l+1})} - Q^0_{(i_l,i_{l+1})}). \end{aligned}$$

Therefore,

$$\begin{aligned} \lim_{n\to\infty} \frac{D^n_{i_L}}{n} &= \lim_{n\to\infty} \Big[\frac{A^n_{m(k)}}{n} - \frac{Q^n_{(0,i_1)} - Q^0_{(0,i_1)}}{n} - \sum_{l=1}^{L-1} \frac{(Q^n_{(i_l,i_{l+1})} - Q^0_{(i_l,i_{l+1})})}{n} \Big] \\ &= \lambda_m = \nu_k, \ a.s. \end{aligned}$$

Now since $\lim_{n\to\infty} \frac{D^n_{k'}}{n} = \nu_k$ *a.s.*, by Lemma 5.5, (k', k) is rate stable. This completes the proof of the induction step and as a result the proof of Theorem 2.2.

5.4. Proof of Theorem 3.2

Let D^n_k denote the cumulative number of tasks that have departed queue k by and including time n: $D^n_k = \sum_{n'=1}^{n} d^{n'}_k$. Define s^n_k to be a random variable denoting the virtual service that queue k receives at time n, whether the queue has been empty or not. s^n_k is a Bernoulli random variable with parameter $\mu_k p^n_k$. Note that $d^n_k = s^n_k \mathbf{1}_{Q^n_k > 0}$. Define the cumulative service that queue k has received up to time n to be $S^n_k = \sum_{n'=1}^{n} s^n_k$.

Lemma 5.6. *The following equality holds:*

$$\lim_{n\to\infty} \frac{\sum_{n'=1}^{n} s_k^{n'} \mathbf{1}_{k\to k'}^{n'}}{n} = \nu_k r_{kk'}, \ a.s., \ \forall k, k'. \tag{61}$$

Proof. First note that the sequence of random variables $\mathbf{1}_{k\to k'}^{n'}$ is i.i.d. Bernoulli-distributed with parameter $r_{kk'}$, and independent of the sequence $s_k^{n'}$. Now by Theorem 3.1, the sequence $\mu_k p_k^{n'}$ converges to ν_k almost surely. Thus, in this probability-1 event, for all $\epsilon_4 > 0$, there exists $n_0(\epsilon_4)$ such that $\|\mu_k p_k^{n'} - \nu_k\| \leq \epsilon_4$ for all $n' > n_0(\epsilon_4)$.

Let w_k^n be an i.i.d Bernoulli process of parameter $(\nu_k - \epsilon_4) r_{kk'}$. We couple the processes $s_k^n \mathbf{1}_{k\to k'}^n$ and w_k^n as follows. If $s_k^n = 0$, then $w_k^n = 0$. If $s_k^n = 1$, then $w_k^n = \mathbf{1}_{k\to k'}^n$ with probability $\frac{\nu_k - \epsilon_4}{\mu_k p_k^n}$, and $w_k^n = 0$ with probability $1 - \frac{\nu_k - \epsilon_4}{\mu_k p_k^n}$. w_k^n is still marginally i.i.d. Bernoulli process of parameter $(\nu_k - \epsilon_4) r_{kk'}$. Then,

$$\liminf_{n\to\infty} \frac{\sum_{n'=1}^{n} s_k^{n'} \mathbf{1}_{k\to k'}^{n'}}{n} \geq \liminf_{n\to\infty} \frac{\sum_{n'=n_0(\epsilon_4)+1}^{n} w_k^{n'}}{n} = (\nu_k - \epsilon_4) r_{kk'} \ a.s.$$

Now we couple the processes $s_k^n \mathbf{1}_{k\to k'}^n$ and v_k^n, where v_k^n is an i.i.d Bernoulli process of parameter $(\nu_k - \epsilon_4) r_{kk'}$. If $s_k^n = 1$, then $v_k^n = \mathbf{1}_{k\to k'}^n$. If $s_k^n = 0$, then $v_k^n = 0$ with probability $\frac{1-(\nu_k+\epsilon_4)}{1-\mu_k p_k^n}$, and $v_k^n = \mathbf{1}_{k\to k'}^n$ with probability $1 - \frac{1-(\nu_k+\epsilon_4)}{1-\mu_k p_k^n}$. v_k^n is still marginally i.i.d. Bernoulli process of parameter $(\nu_k + \epsilon_4) r_{kk'}$. Then,

$$\limsup_{n\to\infty} \frac{\sum_{n'=1}^{n} s_k^{n'} \mathbf{1}_{k\to k'}^{n'}}{n} \leq \limsup_{n\to\infty} \frac{\sum_{n'=n_0(\epsilon_4)+1}^{n} v_k^{n'}}{n} = (\nu_k + \epsilon_4) r_{kk'} \ a.s.$$

The proof is complete by letting $\epsilon_4 \to 0$.

Now we are ready to complete the proof of the theorem. Observe that

$$\begin{aligned} Q_k^n &= Q_k^0 + A_k^n + \sum_{n'=1}^{n} \sum_{k'=1}^{K} d_{k'}^{n'} \mathbf{1}_{k'\to k}^{n'} - D_k^n \\ &\leq Q_k^0 + A_k^n + \sum_{n'=1}^{n} \sum_{k'=1}^{K} s_{k'}^{n'} \mathbf{1}_{k'\to k}^{n'} - D_k^n. \end{aligned}$$

So it is enough to show that

$$\lim_{n\to\infty} \frac{A_k^n + \sum_{n'=1}^{n} \sum_{k'=1}^{K} s_{k'}^{n'} \mathbf{1}_{k'\to k}^{n'} - D_k^n}{n} = 0.$$

First, we show that

$$\liminf_{n\to\infty} \frac{A_k^n + \sum_{n'=1}^{n} \sum_{k'=1}^{K} s_{k'}^{n'} \mathbf{1}_{k'\to k}^{n'} - D_k^n}{n} = 0, \ a.s.$$

In contrary suppose that in a realization,

$$\liminf_{n\to\infty} \frac{A_k^n + \sum_{n'=1}^{n}\sum_{k'=1}^{K} s_{k'}^{n'}\mathbf{1}_{k'\to k}^{n'} - D_k^n}{n} > \epsilon_5,$$

for some $\epsilon_5 > 0$. Then, using Lemma 5.6 and the fact that $\lim_{n\to\infty} \frac{A_k^n}{n} = \lambda_k$, we have

$$\limsup_{n\to\infty} \frac{D_k^n}{n} < \lambda_k + \sum_{k'=1}^{K} \nu_{k'} r_{k'k} - \epsilon_5 = \nu_k - \epsilon_5. \tag{62}$$

On the other hand, $\liminf_{n\to\infty} \frac{Q_k^n}{n} > \epsilon_5$ implies that there exists $n_0(\epsilon_5) > \frac{1}{\epsilon_5}$ such that for all $n > n_0(\epsilon_5)$, $Q_k^n \geq \epsilon_5 n$, or in words, the queue is non-empty after $n_0(\epsilon_5)$. Thus, $s_k^{n'} = d_k^{n'}$ for $n' > n_0$. Therefore,

$$\limsup_{n\to\infty} \frac{S_k^n}{n} \leq \limsup_{n\to\infty} \frac{n_0 + D_k^n}{n} = \limsup_{n\to\infty} \frac{D_k^n}{n}.$$

Now by Lemma 5.4, $\limsup_{n\to\infty} \frac{S_k^n}{n} = \nu_k \leq \limsup_{n\to\infty} \frac{D_k^n}{n}$ which contradicts (62).[6] Since this holds for any $\epsilon_5 > 0$, we conclude that

$$\liminf_{n\to\infty} \frac{Q_k^n}{n} = 0, \ a.s. \tag{63}$$

Second, we show that

$$\limsup_{n\to\infty} \frac{A_k^n + \sum_{n'=1}^{n}\sum_{k'=1}^{K} s_{k'}^{n'}\mathbf{1}_{k'\to k}^{n'} - D_k^n}{n} = 0, a.s.$$

Suppose that in a realization

$$\limsup_{n\to\infty} \frac{A_k^n + \sum_{n'=1}^{n}\sum_{k'=1}^{K} s_{k'}^{n'}\mathbf{1}_{k'\to k}^{n'} - D_k^n}{n} > 2\epsilon_6,$$

for some $\epsilon_6 > 0$. This implies that in this realization, $Q_k^n > 2\epsilon_6 n$ happens infinitely often and

$$\limsup_{n\to\infty} \frac{A_k^n + \sum_{n'=1}^{n}\sum_{k'=1}^{K} s_{k'}^{n'}\mathbf{1}_{k'\to k}^{n'} - D_k^n}{n} > 2\epsilon_6$$

in that realization. Moreover, by (63), for any $\epsilon_6 > 0$, $Q_k^n < \epsilon_6 n$ happens infinitely often with probability 1. Let $N_2 \geq \frac{2}{\epsilon_6}$. Then, there exist $N_2 \leq n_3 < n_4$ such that $Q_k^{n_3} \leq \epsilon_6 n_3$ and $Q_k^{n_4} \geq 2\epsilon_6 n_4$ and queue k is non-empty between times n_3 and n_4. Define a process

$$\tilde{B}_k^n = A_k^n + \sum_{n'=1}^{n}\sum_{k'=1}^{K} s_{k'}^{n'}\mathbf{1}_{k'\to k}^{n'} - S_k^n.$$

[6]The lemma is also valid for the flexible queueing network, and the proof does not change.

Then,

$$\tilde{B}_k^{n_4} = \tilde{B}_k^{n_3} + [\tilde{B}_k^{n_4} - \tilde{B}_k^{n_3}] \tag{64}$$

$$\geq \tilde{B}_k^{n_3} + Q_k^{n_4} - Q_k^{n_3} \tag{65}$$

$$\geq \tilde{B}_k^{n_3} + 2\epsilon_6 n_4 - \epsilon_6 n_3. \tag{66}$$

(65) is due to the following.

$$\begin{aligned}
&\tilde{B}_k^{n_4} - \tilde{B}_k^{n_3} \\
&= A_k^{n_4} - A_k^{n_3} + \sum_{n'=n_3+1}^{n_4} \sum_{k'=1}^{K} s_{k'}^{n'} \mathbf{1}_{k' \to k}^{n'} - (S_k^{n_4} - S_k^{n_3}) \\
&\geq A_k^{n_4} - A_k^{n_3} + \sum_{n'=n_3+1}^{n_4} \sum_{k'=1}^{K} d_{k'}^{n'} \mathbf{1}_{k' \to k}^{n'} - (S_k^{n_4} - S_k^{n_3}) \\
&= A_k^{n_4} - A_k^{n_3} + \sum_{n'=n_3+1}^{n_4} \sum_{k'=1}^{K} d_{k'}^{n'} \mathbf{1}_{k' \to k}^{n'} - (D_k^{n_4} - D_k^{n_3}) \qquad (67) \\
&= Q_k^{n_4} - Q_k^{n_3}.
\end{aligned}$$

(67) is true since queue k is non-empty between times n_3 and n_4. Now (66) implies that

$$\frac{\tilde{B}_k^{n_4}}{n_4} - \frac{\tilde{B}_k^{n_3}}{n_3} \geq \frac{\tilde{B}_k^{n_3}}{n_3}(\frac{n_3}{n_4} - 1) + 2\epsilon_6 - \epsilon_6 \frac{n_3}{n_4}.$$

By Lemma 5.4 we know that

$$\lim_{n \to \infty} \frac{\tilde{B}_k^n}{n} = 0 \ a.s.,$$

so one can choose N_2 large enough such that for all $n \geq N_2$, $|\tilde{B}_k^n / n| \leq \epsilon_6/3$. Then,

$$|\frac{\tilde{B}_k^{n_4}}{n_4} - \frac{\tilde{B}_k^{n_3}}{n_3}| \leq \frac{2\epsilon_6}{3}. \tag{68}$$

However, since $\frac{\tilde{B}_k^{n_3}}{n_3} \leq \epsilon_6$ and $\frac{n_3}{n_4} < 1$,

$$\frac{\tilde{B}_k^{n_4}}{n_4} - \frac{\tilde{B}_k^{n_3}}{n_3} \geq \epsilon_6(\frac{n_3}{n_4} - 1) + 2\epsilon_6 - \epsilon_6 \frac{n_3}{n_4} = \epsilon_6,$$

which contradicts (68). Thus,

$$\limsup_{n \to \infty} \frac{A_k^n + \sum_{n'=1}^{n} \sum_{k'=1}^{K} s_{k'}^{n'} \mathbf{1}_{k' \to k}^{n'} - D_k^n}{n} = 0, a.s.,$$

which completes the proof of Theorem 3.2.